


CONNEXIONS



The Interoperability Report

January 1993

Volume 7, No. 1

ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.

In this issue:

Components of OSI: International Standardized Profiles.....	2
The Message Send Protocol...	16
Inside an IETF Plenary.....	18
NREN Politics Thicken.....	23
SLIP at the University of Minnesota.....	24
Announcements.....	27
Book Reviews.....	30

From the Editor

Our series *Components of OSI* continues this month with an examination of the *International Standardized Profiles* (ISPs). Since the OSI standards are full of options, it is necessary to identify preferred combinations of standards and options for various target environments. The resulting specifications are called *functional standards*, and are being produced by groups such as MAP/TOP, SPAG, POSI and government organizations through the various GOSIP programs. ISPs attempt to manage the development of functional standards in an orderly and consistent manner. The article is written by William Stallings, and is adapted from his new book *Networking Standards: A Guide to OSI, ISDN, LAN, and MAN Standards*, which will be published by Addison-Wesley this month.

Billy Barron gives an overview of the *Message Send Protocol* which is used to send "one-liner" messages to other users on a local or remote host. This functionality has existed in many operating systems over the years, but has not been implemented in a consistent network-wide fashion until recently.

Last November, the *Internet Engineering Task Force* (IETF) met in Washington, D.C. It was perhaps fitting that this meeting should take place so close to Capitol Hill, as much of the discussion centered around a new "constitution" and the "consent of the governed" with respect to the relationship between the Internet Society, Internet Architecture Board (IAB), Internet Engineering Steering Group (IESG), and the IETF itself. This effort, known in the community as "POISED," will be discussed in a future issue. This time, we will take you inside a "regular" IETF plenary meeting and let you know what to expect if you've never attended such an event. The author of the article is Gary Malkin who has been active in helping newcomers understand the workings of the IETF.

NREN updates have become almost a regular feature in *ConneXions*. We asked Mike Roberts to once again give us a status report as the country transitions to a new administration.

Finally, another case study from the University of Minnesota, this time on their use of the *Serial Line IP* (SLIP) protocol. SLIP continues to be an attractive low-cost way of expanding IP connectivity to home PCs and the like. The article is by Craig Finseth.

Included with this issue of *ConneXions* is the table of contents for all back issues from "the beginning" (March 1987) through December of 1992. All back issues are available for purchase, as complete volumes (this includes a special *ConneXions* binder) or single copies.

ConneXions is published monthly by Interop Company, 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779. Toll-free: 1-800-INTEROP.
E-mail: connexions@interop.com.

Copyright © 1993 by Interop Company.
Quotation with attribution encouraged.

ConneXions—The Interoperability Report and the *ConneXions* logo are registered trademarks of Interop Company.

ISSN 0894-5926

Components of OSI: International Standardized Profiles

by William Stallings, Comp-Comm Consulting

Introduction

Although the OSI model was first defined 15 years ago and has been an international standard for over a decade, OSI implementations have been slow to come to market. A key reason for this has to do with the magnitude of the task of transitioning from another communication suite to OSI. With well-established proprietary architectures, such as SNA, and a widely used multivendor architecture, namely the TCP/IP protocol suite, it was inevitable that there would be inertia to overcome. But the pace of development and adoption of OSI has also been hampered by some practical problems relating to the standards themselves:

1. Most of the ISO and CCITT standards relating to OSI are based on paper designs, usually developed by committee. That is, a functional requirement, such as for a connection-oriented transport protocol, is developed and then a specification for the protocol is drawn up. It is only after the specification has solidified that serious implementation efforts begin. It is possible, indeed likely, that some subtle logical flaws exist in such specifications that only show up after field experience with the implementation.
2. The protocol specifications are essentially prose descriptions. Although they are often accompanied by some formal description, such as finite state machine definitions, the protocol specifications are not amenable to mathematical proof nor is it possible to mathematically verify an implementation. Thus, there is opportunity for ambiguity and the risk that two different implementations will not interoperate.
3. A standard typically contains a number of options and parameters with undefined value ranges. This allows a single protocol specification to satisfy a range of requirements that are similar but not identical in a consistent manner. However, two implementations that make different choices for options and parameters may not be able to interoperate.

Problem 1 is dealt with by experience. Many of the ISO and CCITT standards have now been around for a number of years and, as the body of experience grows, so does the confidence in the correctness of the specifications. For newer protocols, the amount of experience is less but growing.

Problem 2 and, to a certain extent problem 3, are dealt with by the OSI conformance testing process. [See December 1992 *ConneXions*]. To the extent that conformance assessment is robust and correct, we can have a degree of confidence that an implementation conforms to the corresponding protocol specification and that the options and parameters chosen for the implementation are clearly documented.

Problem 3 is more significant than it might first appear. With virtually all of the protocol standards, there is a substantial number of options and adjustable parameters. As an example, consider ISO 9542, the connectionless end-system-to-intermediate-system routing protocol. This is one of the simplest of the OSI protocols; the actual specification of the protocol takes just ten pages; however, the *Protocol Implementation Conformance Statement* (PICS) proforma, which defines the options and parameters for the protocol, takes an additional five pages. The PICS proformas for more complex protocols are correspondingly longer.

Functional standards

Problem 3 becomes even more significant when we consider that no user is interested in a single protocol at a single layer of the OSI architecture. Rather, the user is interested a full 7-layer implementation that will support one or more distributed applications (e.g., file transfer). In order to specify the OSI requirements to satisfy a particular application, it is necessary to reference a set of standards and specify, for each standard referenced, the valid options and parameter settings within that standard that are needed to achieve the required function. Considering that collectively, across 7 layers, there will be a large number of options and parameters, it is not unlikely that two different implementations that claim to support a given application may make different choices and not be interoperable.

To address this problem, a number of organizations have created documents that identify preferred combinations of standards and options for various application areas. The resulting documents have been termed *functional standards*. Two examples of this are the MAP/TOP specification, which deals with office and factory applications, and the GOSIP specification, which deals with government requirements. Other notable examples include the *Standards Promotion and Application Group* (SPAG) in Europe, operated by the European Commission, and the *Promoting Conference for OSI* (POSI) in Japan. All of these regional and special-purpose efforts represent an improvement over an ad hoc attempt to specify application-based requirements. However, these same efforts run the risk of fragmenting what should be a global market and adding further confusion to that already facing the customer and end user.

International Standardized Profiles

In an effort to manage the development of functional standards in an orderly and consistent manner, ISO/IEC JTC1 (*Joint Technical Committee 1*) established the *Special Group on Functional Standardization* (SGFS) to bring these regional and special-purpose activities together into a single program to provide functional standards of global significance. To achieve its purpose, SGFS created a new type of publication called the *International Standardized Profile* (ISP). In mid-1990, ISO published TR 10000, "Framework and Taxonomy of International Standardized Profiles," and since then has published a number of ISPs. A key objective of this effort is to minimize the confusion on the part of users and suppliers by developing a common classification scheme, document scope and document style for functional standards. Table 1 lists the currently published ISPs.

TR 10000-1 defines the concept of profiles, and the way in which they are documented in International Standardized Profiles (ISPs). It also sets out the nature and content of the documentation required for an ISP. The document lists the following purposes for defining profiles:

- Identify base standards, together with appropriate classes, subsets, options, and parameters that are necessary to accomplish identified functions.
- Provide a system of referencing the various uses of base standards that is meaningful to both users and suppliers.
- Provide a means to enhance the availability for procurement of consistent implementations of functionally-defined groups of base standards, which are expected to be the major components of real application systems.
- Promote uniformity in the development of conformance tests for systems that implement the functions associated with the profiles

International Standardized Profiles *(continued)*

ISO TR 10000-1	International Standardized Profiles, Part 1: Framework
ISO TR 10000-2	Part 2: Taxonomy of Profiles
ISP 10607-1	ISP AFTnn — FTAM, Part 1: Specification of ACSE, Presentation, and Session Protocols for the use by FTAM
ISP 10607-2	Part 2: Definition of Document Types, Constraint Sets, and Syntaxes
ISP 10607-2 DAD 1	Addendum 1: Additional Definitions
ISP 10607-3	Part 3: AFT11 — Simple File Transfer Service
ISP 10607-4	Part 4: AFT12 — Positional File Transfer Service (Flat)
DISP 10607-5	Part 5: AFT22 — Positional File Access Service (Flat)
DISP 10607-6	Part 6: AFT3 — File Management Service
DISP 10608-1	Connection-mode Transport Service over Connectionless Network Service, Part 1: General Overview and Subnetwork-Independent Requirements
DISP 10608-2	Part 2: TA51 Profile Including Subnetwork-Dependent Requirements for CSMA/CD LANs
DISP 10608-5	Part 5: TA1111/TA1121 Profiles Including Subnetwork-Dependent Requirements for X.25 Packet Switched Data Networks Using Switched Virtual Circuits
DISP 10609-1	ISPs TB, TC, TD, and TE — Connection-mode Transport Service over Connection-mode Network Service, Part 1: Subnetwork-type Independent Requirements for Group TB
DISP 10609-2	Part 2: Subnetwork-type Independent Requirements for Group TC
DISP 10609-3	Part 3: Subnetwork-type Independent Requirements for Group TD
DISP 10609-4	Part 4: Subnetwork-type Independent Requirements for Group TE
DISP 10609-5	Part 5: Definition of Profile TB 1111/TB 1121
DISP 10609-6	Part 6: Definition of Profile TC 1111/TC 1121
DISP 10609-7	Part 7: Definition of Profile TD 1111/TD 1121
DISP 10609-8	Part 8: Definition of Profile TE 1111/TE 1121
DISP 10609-9	Part 9: Subnetwork-type Dependent Requirements for Network Layer, Data Link Layer, and Physical Layer Concerning Permanent Access to a Packet Switched Data Network Using Virtual Call
DISP 10610-1	ISP FOD11 — Office Document Format — Simple Document Structure — Character Content Architecture Only — Part 1: Document Application Profile
DISP 11181-1	ISP FOD26 — Office Document Format — Enhanced Document Structure — Character, Raster Graphics and Geometric Graphics Content Architectures — Part 1: Document Application Profile
DISP 11182-1	ISP FOD36 — Office Document Format — Extended Document Structure — Character, Raster Graphics and Geometric Graphics Content Architectures — Part 1: Document Application Profile
DISP 1183-1	ISP AOMnn OSI Management — Management Communication Protocols — Part 1: Specification of ACSE, Presentation and Session Protocols for the Use by ROSE and CMISE
DISP 1183-2	Part 2: AOM12 — Enhanced Management Communications
DISP 1183-3	Part 3: AOM11 — Basic Management Communications

Table 1: International Standardized Profiles

TRANSPORT PROFILES: SUBNETWORKS

TRANSPORT PROFILES: SUBNETWORKS		TRANSPORT PROFILES: TRANSPORT GROUPS		APPLICATION PROFILES: MESSAGE HANDLING (AMH)		APPLICATION PROFILES: REMOTE DATABASE ACCESS (ARD)	
1	PACKET SWITCHED DATA NETWORKS (PSDN)	TA	CO-TS over CL-NS ¹	1	COMMON FACILITIES	1	Substructure to be studied
11	Permanent Access to a PSDN	TB	CO-TS over CO-NS with mandatory protocol classes 0, 2, and 4	11	MTA and MTS ⁵	11	Substructure to be studied
111	Virtual call	TC	CO-TS over CO-NS with mandatory protocol classes 0 and 2	12	UA to MS (P7)	12	APPLICATION PROFILES: OSI MANAGEMENT (AOM)
1111	Permanent virtual circuit	TD	CO-TS over CO-NS with mandatory protocol class 0	13	UA or MS to MTA (P3)	13	Substructure to be studied
1112	Digital data circuit/CSDN leased line	TE	CO-TS over CO-NS with mandatory protocol class 2	2	IPMS	2	Substructure to be studied
112	Virtual call	UA	CL-TS over CL-NS	21	IPM end system to IPM end system (P2 over P1)	21	Substructure to be studied
1121	Permanent virtual circuit	UB	CL-TS over CO-NS	22	IPM UA to IPM MS (P2 over P7)	22	APPLICATION PROFILES: DIRECTORY (ADI)
1122	ISDN B channel, semi-permanent			23	IPM UA or IPM MS to MTA (P2 over P3)	23	
113	Virtual call			24	IPM end system to IPM end system (P2-1984 over P1-1984)	24	
1131	Permanent virtual circuit						
1132	Switched access to a PSDN						
12	PSDN case						
121	Virtual call						
1211	CSDN case						
122	Virtual call						
1221	ISDN B channel case						
123	Virtual call						
1231	Virtual call						
2	DIGITAL DATA CIRCUIT						
21	Leased (permanent) service						
22	Dial-up (CSDN)						
3	ANALOG TELEPHONE CIRCUIT						
31	Leased (permanent) service						
32	Dial-up (PSTN)						
4	INTEGRATED SERVICES DIGITAL NETWORK (ISDN)						
41	Semi-permanent service						
411	B channel						
4111	X.25 DTE to DTE operation						
42	Circuit mode service						
421	B channel						
4211	X.25 DTE to DTE operation						
43	Packet mode service						
431	D channel access						
4311	Virtual call						
4312	Permanent virtual circuit						
432	B channel semi-permanent access						
4321	Virtual call						
4322	Permanent virtual circuit						
433	B channel demand access						
4331	Virtual call						
5	LOCAL AREA NETWORKS						
51	CSMA/CD						
52	Token bus						
53	Token ring						
54	FDDI						

Table 2: Taxonomy of Profiles

Relationship of Profile
to Base Standards

International Standardized Profiles (continued)

The relationship between a profile and the base standards to which it refers can be defined in terms of three concepts: *selection*, *limitation*, and *conformance*.

A profile is a *selection* of base standards that, in combination, can be used to provide a given function in a given environment. For each base standard selected, a choice is made of permitted options within that standard. In addition, suitable values for parameters which are left unspecified in the base standard are provided. In some cases, a parameter is completely unspecified in the base standard. In other cases, a default or recommended value is provided but other values may be selected for implementation.

A profile is also a *limitation* of the base standards for the given application. The choice of options and ranges of values may be restricted so as to maximize the probability of interworking between different systems that conform to the profile and that make selections among the options and parameters ranges remaining in the profile. Of course, the choices made in the profile must not contradict the base standards. In particular, if certain combinations of options and/or parameter values are forbidden in the base standard, then they must also be forbidden in the profile. If the development of the profile indicates the need to modify or add to the requirements specified in the base standard, these modifications must be made in the base standard and not simply incorporated in the profile.

A profile may contain *conformance* requirements that are more specific and limited in scope than those of the base standards to which the profile refers. While conformance to a profile always implies conformance to the set of base standards that the profile references, the opposite is not always true.

Framework and
taxonomy

The task of specifying interoperable protocols across seven layers for a given application can be made more manageable by organizing the OSI architecture into four relatively independent components, as illustrated in Figure 1. For a given application, a functional standard, or profile, is needed for each component: application, transport protocol, internetwork protocol, and subnetwork.

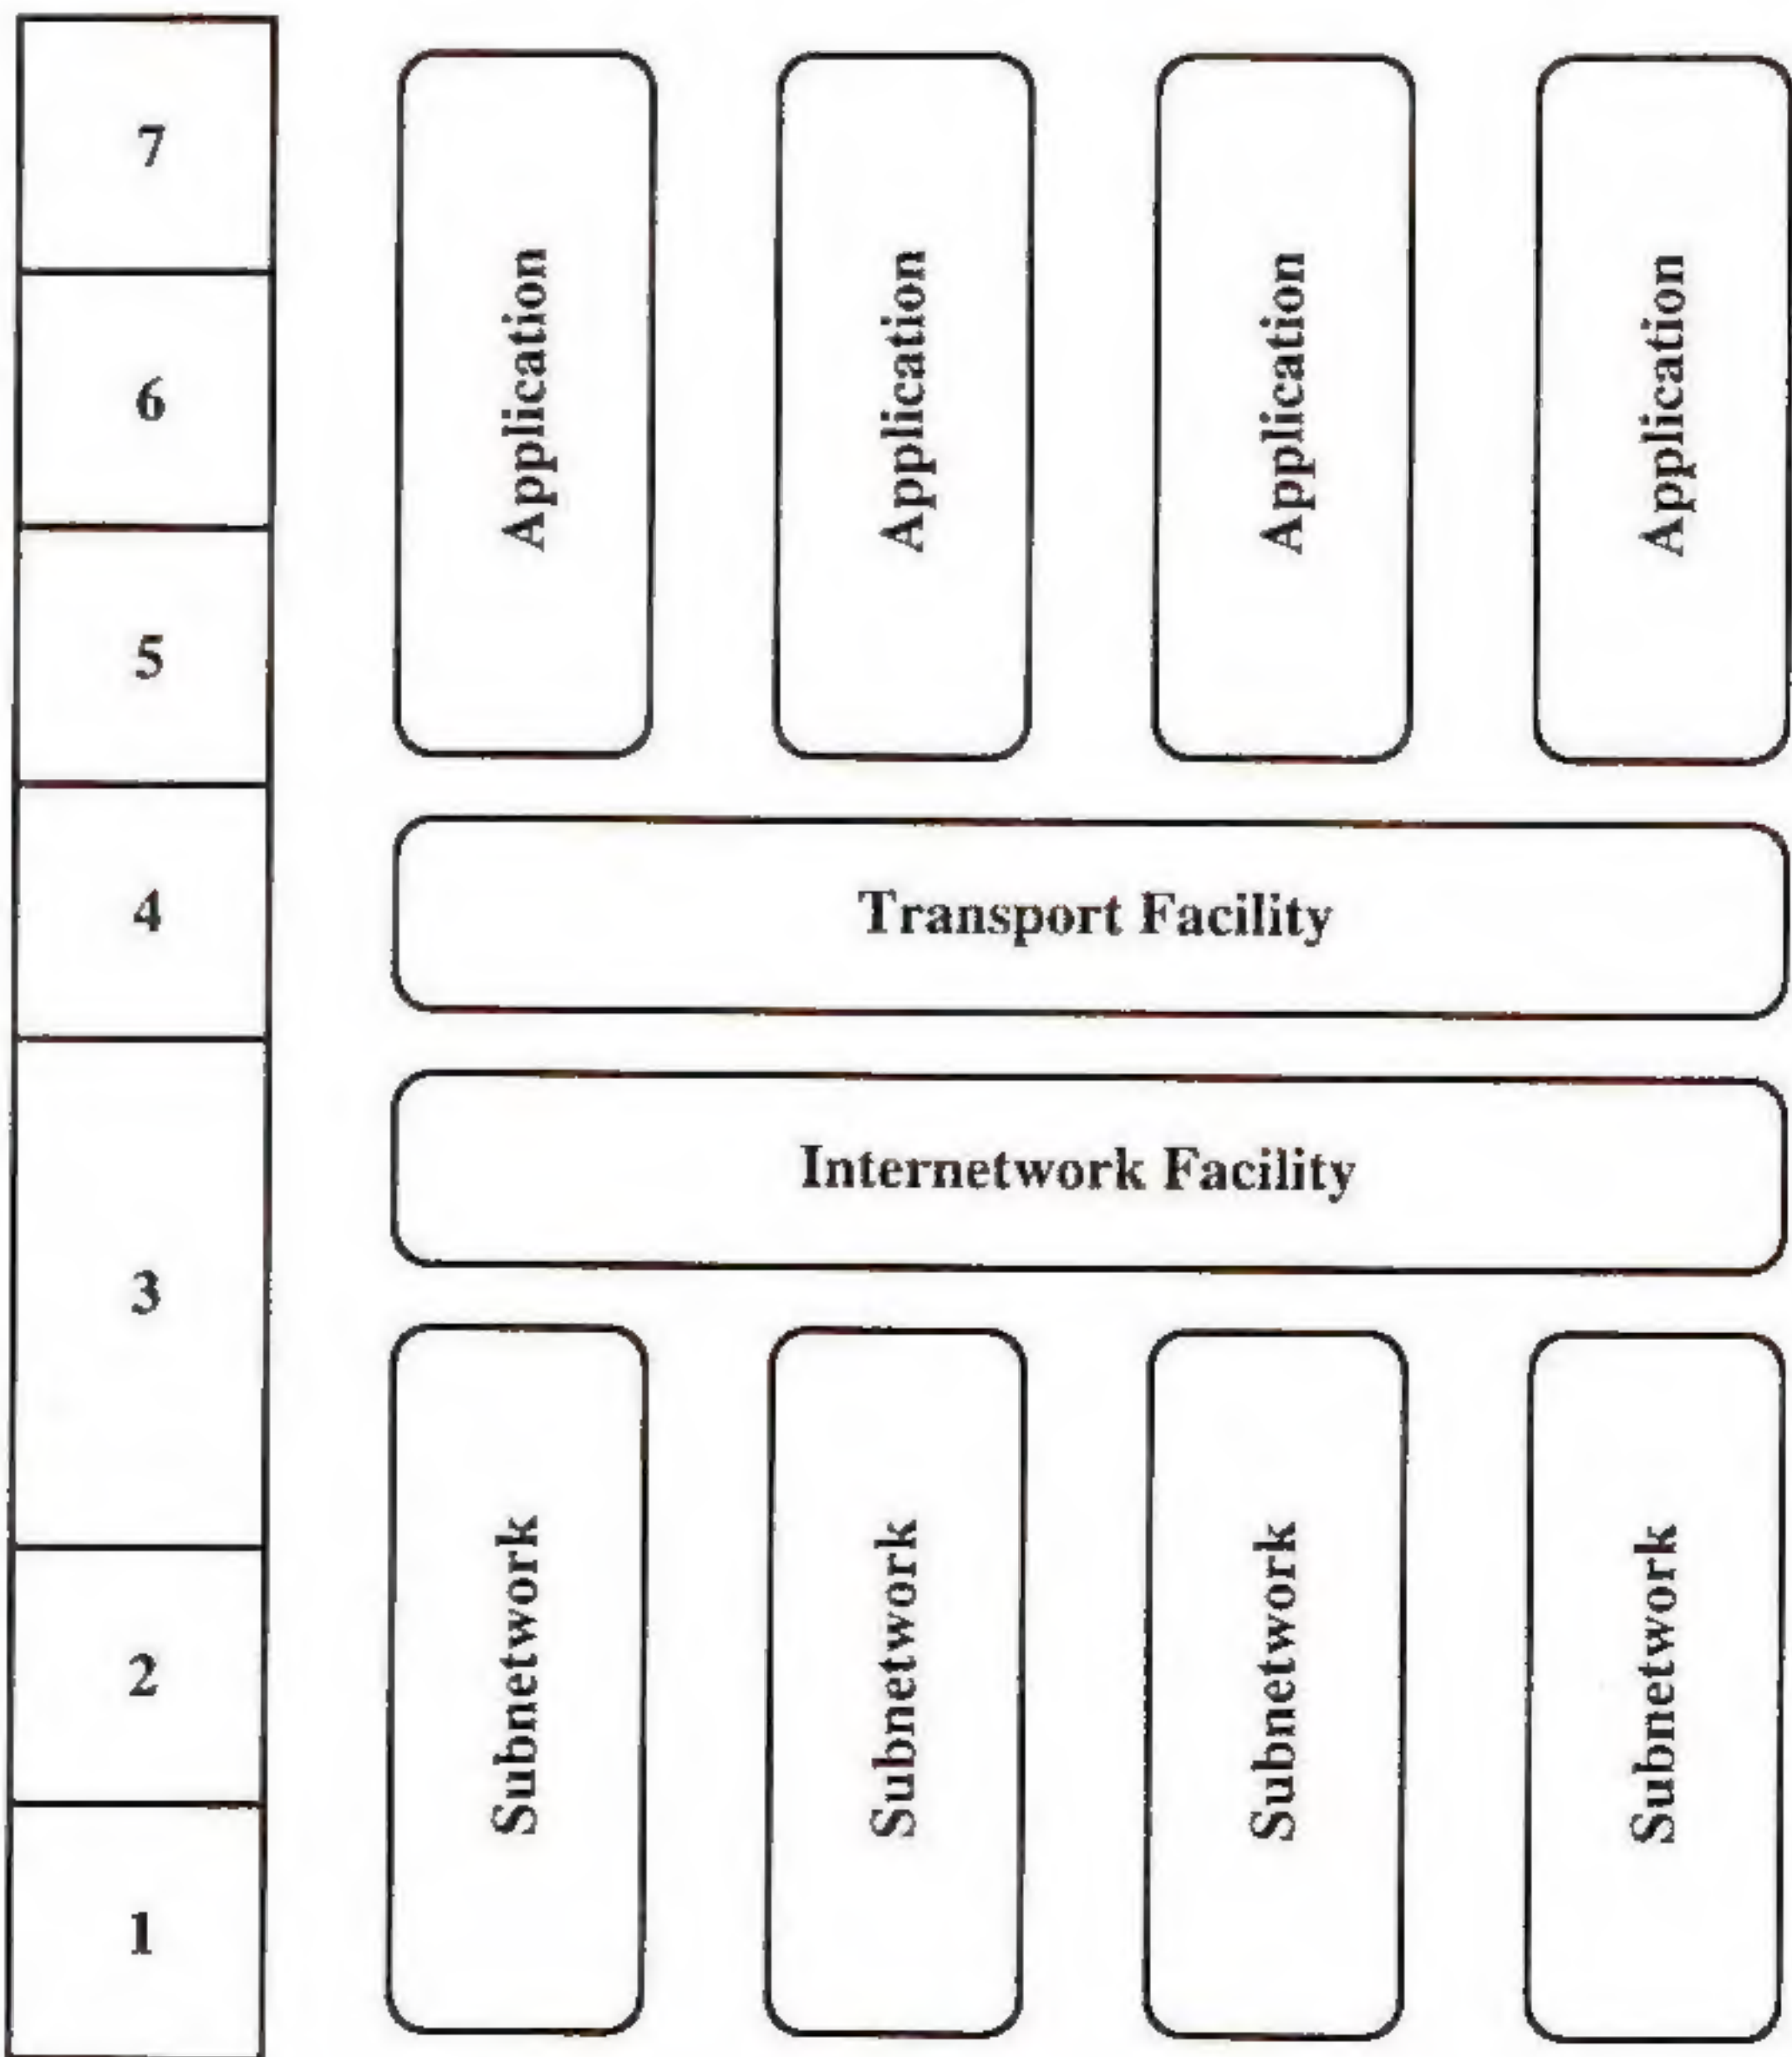


Figure 1: Key Components for Interoperability

The choices that need to be made for each are best described in the following order:

- *Application Suite*: A given application suite (This is not an official OSI term, but is a useful concept) will consist of a set of *Application Service Elements* (ASEs) plus the presentation functional units and session functional units specifically needed to support the application.
- *Subnetwork*: A subnetwork or subnetworks interconnect the distributed systems. The set of subnetworks may be imposed; that is, a new application and perhaps new systems are added to an existing network facility. Alternatively, the subnetwork or subnetworks may be chosen for one or more applications. In this latter case, the choice is likely to be based on network technology, capacity requirements, and a variety of other factors that do not necessarily reflect the particular characteristics of the application.
- *Internetworking Facility*: Given a collection of subnetworks, an internetworking protocol and the associated ES-IS and IS-IS routing protocols must be selected. The choice of either connection-mode or connectionless-mode will depend on the nature of the subnetworks.
- *Transport Facility*: Either the connection-mode or connectionless-mode transport service is chosen to match the mode of the application. Then, the appropriate class of transport protocol is chosen to match the characteristics of the internetworking facility and the underlying subnetworks.

For all four components, the appropriate options and parameter settings to support the given application must be selected.

Profile reuse

The advantage of viewing the profiling process in this fashion is the potential for the reuse of profiles. For the application suite, a different profile will be defined for each separate application. A small number of general-purpose profiles can be defined for the transport facility, internetworking facility, and subnetworks. A total profile for an application then consists of a profile for the application suite together with a selection of the appropriate general-purpose profiles for each of the other three components.

One final point concerns the relationship between connection-mode and connectionless-mode: a conversion from one mode to another can occur from the transport service to the network service and from the network service to the subnetwork capability. Both of these types of conversions need to be considered in developing a full profile to support an application.

TR 10000-2 defines a system of labels for the profiles, which provides a structure and classification within which the profiles will fit. These labels reflect the applicability of the profile and its constraints. This system of labels is called the *taxonomy*. There are two reasons for developing a taxonomy:

- It provides a technical framework for the development of ISPs. The taxonomy reflects the structure of real systems (as suggested by Figure 1) and hence guides the development of ISPs in a way that is most useful for actually defining practical implementations. The framework thus serves to coordinate the work of the various groups interested in developing ISPs.
- It helps users identify the particular profile or profiles that address their requirements.

International Standardized Profiles *(continued)*

The taxonomy provides a hierarchical structure. At a top level, profiles are divided into *classes*, each class representing a category of functionality of reasonable independence from other classes. Each class is further subdivided recursively. Each leaf of the tree represents a profile, and the path to that leaf is the identifier of the profile. Table 2 is the taxonomy at its current stage of development.

Profile classes

Six profile classes are defined in TR 10000-2:

- F: Interchange format and representation profiles
- A: Application profiles requiring the connection-mode transport service (COTS)
- B: Application profiles requiring the connectionless-mode transport service (CLTS)
- T: Transport profiles that provide the connection-mode transport service (COTS)
- U: Transport profiles that provide the connectionless-mode transport service (CLTS)
- R: Relay profiles that define relay functions between T or U profiles

Thus, the profile classes divide OSI functionality into profiles that deal with the structuring and coding of information, profiles that deal with the application and its communication needs, and profiles that are concerned with the use of a network technology to achieve the interconnection requirements. To provide a full implementation of a given function or application, a combination of profiles will be needed. Figure 2 illustrates the relationships among classes and indicates the allowable combinations. A and T profiles can be combined since the T profile provides the service required by the A profile. Similarly, B and U profiles can be combined. Conversely, combinations of A and U profiles or B and T profiles are not possible.

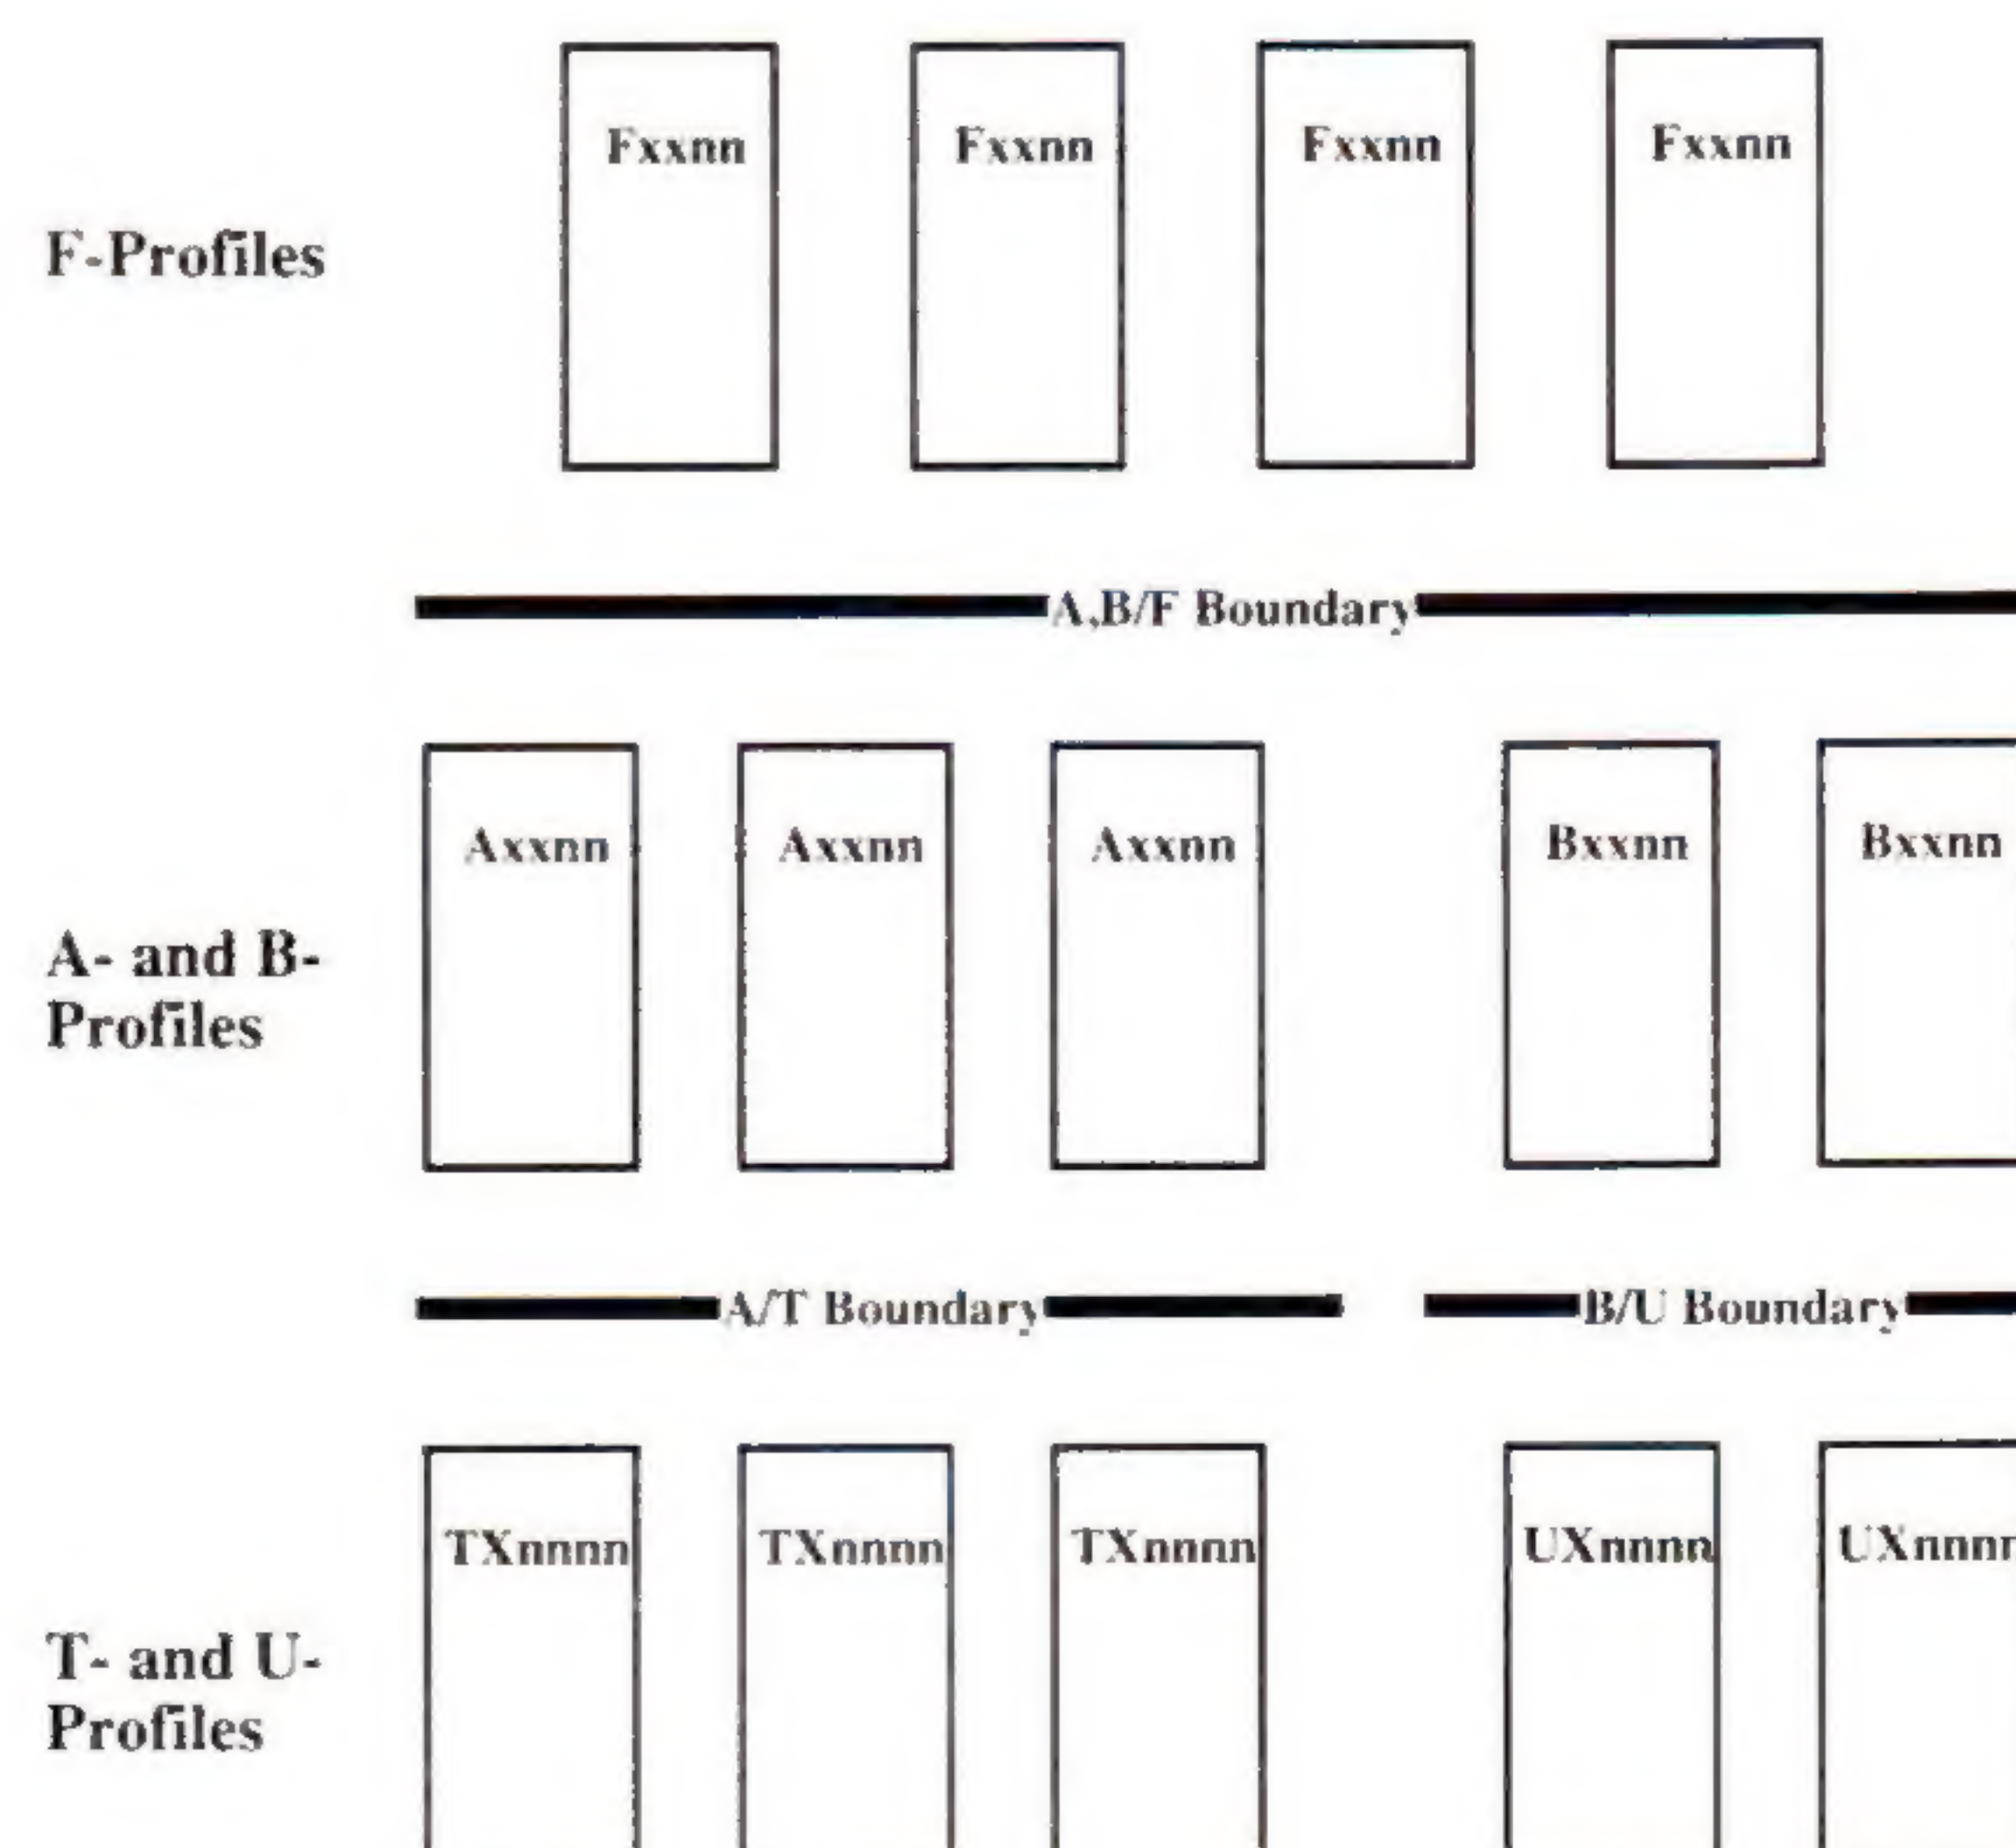


Figure 2: Relationship between Profiles in the OSI Taxonomy

Since there is no single service definition that characterizes the A/F and B/F boundaries, there is no restriction on combining F profiles with A or B profiles at this level of the taxonomy. However, restrictions may exist in the base standards or in the profiles to limit the number of possible combinations.

Transport profiles

Transport profiles specify how the two modes of transport service are provided over the two modes of network service and over specific sub-network types. The transport profiles are subdivided into seven *groups*, five groups for class T and two groups for class U. The group concept is defined such that the profiles within a single group are compatible in the sense that a system implementing one profile from a group can interwork with another system implementing another profile from the same group.

There are three factors that guarantee interworking between two transport profiles:

- They provide the same mode of transport service.
- They utilize the same mode of network service.
- In the case of COTS over CONS, different protocol classes (0, 1, 2, 3, 4) may be implemented. As a result, it is possible for two implementations of COTS over CONS to be unable to agree on a protocol class. Thus, the requirement is for support of the same subset of transport protocol classes.

With these requirements satisfied, a group can contain profiles which correspond to different subnetwork technologies; interworking between systems in the group is made possible by LAN bridges and/or network-layer relays (intermediate systems). Based on the above line of reasoning the following groups are identified:

- TA: COTS over CLNS (transport class 4 is mandatory)
- TB: COTS over CONS (transport classes 0, 2, & 4 are mandatory)
- TC: COTS over CONS (transport classes 0 & 2 are mandatory)
- TD: COTS over CONS (transport class 0 is mandatory)
- TE: COTS over CONS (transport class 2 is mandatory)
- UA: CLTS over CLNS
- UB: CLTS over CONS

Within each group, a number of subnetworks are identified. Table 2 lists all of the subnetworks so far identified, together with the restrictions that apply. The identifier for a particular transport profile indicates the class and group with two leading capital letters, followed by one or more digits to indicate the subnetwork (Table 3).

Transport Profiles: <i>CXabcd</i>	
where:	
C	= transport class designator (T or U)
X	= a letter that identifies the group within the class
abcd	= numerical identifier of the subnetwork type supported in this profile
Relay Profiles: <i>RXp.q</i>	
where	
R	= relay function
X	= relay type identifier
p, q	= subnetwork identifiers
Application Profiles: <i>CXYabc</i>	
where	
C	= application class designator (A or B)
XY	= two letters that identify the primary subdivision, taken from the main categories of application functions and OSI management
abc	= numerical identifier for the member(s) of the subdivision
Interchange Format and Representation Profiles: <i>FXYabc</i>	
where	
F	= interchange format class designator
XY	= two letters that identify the primary subdivision
abc	= numerical identifier for the member(s) of the subdivision

Table 3: Summary of Profile Label Formats

continued on next page

International Standardized Profiles (*continued*)

A small letter in place of a digit indicates that all members of a subgroup are covered. For example:

Profile TA 51: Connection-mode transport service over connectionless-mode network service supporting a CSMA/CD local network.

Profile TC 112x: Connection-mode transport service over connection-mode network service with mandatory protocol classes 0 and 2, supporting a PSDN with permanent access via a digital data circuit or CSDN leased line.

Interworking

The group concept dictates that interworking is possible between systems using profiles in the same group. It is also possible to consider interworking between systems in different groups, but care must be taken. In some cases, the normal transport protocol class negotiation mechanism will allow the systems to achieve a level of communication via OSI relays. In other cases, a non-OSI relay will be required that performs some sort of protocol conversion.

At a top level, we can say that no interworking is possible between a group in class T and a group in class U, because of the different mode of transport service provided. Tables 4a and b show the interworking possibilities among groups within the same class. These tables provide some indication of the difficulty in providing interworking between profiles in different groups.

(a) Interworking between groups in Class T

Responder in Group	Network Service Module	Initiator in Group				
		TA	TB	TC	TD	TE
TA	CL	full	special 1	special 1	special 1	special 1
TB	CO	special 1	full	full	full	full
TC	CO	special 1	restricted	full	full	full
TD	CO	special 1	restricted	restricted	full	special 2
TE	CO	special 1	restricted	restricted	special 2	full

(b) Interworking between groups in Class U

Responder in Group	Initiator in Group	
	UA	UB
UA	full	special 2
UB	special 2	full

Key to tables:

- full:** Full interworking is possible (one or more OSI relays may be required).
- restricted:** Interworking is restricted in that it may not be possible to use the class of transport protocol preferred by the initiator.
- special 1:** Non-OSI relay required for interworking; special restrictions for interworking exist.
- special 2:** Non-OSI relay required for interworking; interworking between these profile types is not contemplated.

Table 4: Interworking Between Groups

Relay profiles

Relay profiles define the use of standards from OSI layers 1 to 4 to provide relaying functions between OSI transport profiles. At a top level, the following subclasses of relays have been identified:

- RA: Relaying the CLNS
- RB: Relaying the CONS
- RC: X.25 protocol relaying
- RD: Relaying the MAC service using transparent bridging
- RE: Relaying the MAC service using source routing
- RZ: Relaying between CLNS and CONS

A relay links two subnetworks. Hence such subnetworks are used as a next level of distinction for relay profiles. Examples:

Profile RA 51.53: Relays the connectionless-mode network service between a CSMA/CD LAN and a token ring LAN.

Profile RC 22.51: Relays the X.25 protocol between a CSDN and a CSMA/CD LAN.

Application profiles

Application profiles define the use of protocol standards encompassing OSI layers 5 through 7, to provide for the structured transfer of information between end systems. As with transport profiles, application profiles are divided into two classes: class A application profiles are connection-mode and therefore require COTS; class B application profiles are connectionless-mode and therefore require CLTS.

Each application profile is a complete definition of the use of protocol standards from layers 5 through 7, including a selection of application service elements, session functional units, and presentation functional units. Currently, only class A profiles have been defined:

- AFT: File transfer, access, and management (FTAM)
- AMH: Message handling (based on X.400)
- AVT: Virtual terminal
- ATP: Transaction processing
- ARD: Remote database access
- AOM: OSI management
- ADI: Directory

Interchange format and representation profiles

Interchange format and representation profiles define the structure and/or content of the information being interchanged by application profiles. Hence, the main feature that distinguishes them from application profiles is the absence of a transfer function. The following profiles have been defined:

- FOD: Office document format
- FCG: Computer graphics metafile interchange format
- FSG: Standard generalized markup language interchange format
- FDI: Directory data definitions

The ISP document

TR 10000-1 dictates a specific structure for an ISP document. The document begins with several *preliminary elements*, that provide an informal description of the ISP:

- *Foreword:* Indicates the organization or committee that prepared the ISP and indicates whether this ISP replaces or amends a previous ISP.
- *Introduction:* Describes the process used to draft the ISP and the degree of international harmonization that it has received.

The key portion of the document are the *normative elements*, which set out the provisions with which it is necessary to comply in order to be able to claim conformity with the ISP:

- *Scope:* Defines the purpose and subject matter of the ISP, relates the ISP to the taxonomy published in TR 100000-2, and includes the scenario of the profile. This latter is an illustration of the environment within which the ISP is applicable.

International Standardized Profiles (*continued*)

- *Normative References*: A list of normative documents (International Standards, Technical Reports, ISPs, and CCITT Recommendations) to which reference is made in the text.
- *Definitions*: Provides definitions necessary to the understanding of certain terms used in the ISP.
- *Abbreviations*: Provides a list of the symbols and abbreviations used in the ISP.
- *Requirements*: Includes clauses relating to the use made of each of the base standards referenced in the profile definition. This includes, the choice of classes or subsets of the overall specification, the selection of options, and the selection of ranges of parameter values. Although the exact format is not dictated, it shall be in the form of static and dynamic conformance requirements. This section may be quite brief, with most of the details concerning choices made of classes, subsets, options, and ranges of parameter values recorded in an annex.
- *Normative Annexes*: These are integral sections of the ISP that, for reasons of convenience, are not included in the main body of the document.

Finally, the document contains *supplementary elements*, which provide additional information intended to assist the understanding or use of the ISP.

For similar profiles that cover several layers, there will be many cases in which the same set of options from a particular standard is called up by different profiles. To handle this situation efficiently, an ISP may be published as a set of related parts (multi-part ISP) rather than a single document. A multi-part ISP will encompass a number of related profiles.

ISP 10609

One example of this structure is ISP 10609, which covers profiles TB, TC, TD, and TE, all of which provide COTS over CONS. Figure 3 shows how the group structure leads to a modular structure for the definition of the profiles within a group, with references to common elements of text. Parts 1 – 4 of the ISP provide the requirements that are independent of subnetwork type for the four COTS over CONS groups. The only difference between the groups is which transport protocol classes are mandatory. Each group references the connection-mode transport protocol (8073) and service (8072) standards and the connection-mode network service (8348) definition. Since the four groups are so similar, many of their requirements are the same and are therefore independent of group. These group-independent and subnetwork-type-independent requirements are listed only once in Part 1. Parts 2 through 4 reference Part 1 for these requirements.

The next four parts of 10609 correspond one for one with the first four parts and deal with a particular subnetwork type: 111, which is leased-line access via a *public switched telephone network* (PSTN) to a *packet-switched data network* (PSDN). Each part specifies the use of the X.25 protocol stack (layers 1, 2, and 3) to support the OSI connection-mode network service. Thus, each references various physical-layer standards, LAPB (7776), and the X.25 layer 3 (8208 and 8878). None of these parts contain any technical specification, but references other ISP parts. Thus part 5 references Part 1; Part 6 references part 2; and so on. In addition, all four parts reference Part 9, which provides the subnetwork-type-dependent requirements. In this case, these requirements are for support of CONS over an X.25 PSDN.

The use of a multi-part structure reduces the amount of duplicated text and the risk of unintentional differences between profiles. In this case, the number of documents required to specify a single profile is large, but as more profiles are defined, there will be more re-use of modules. As an example, ISO 10608 includes a cross-reference to a part of 10609.

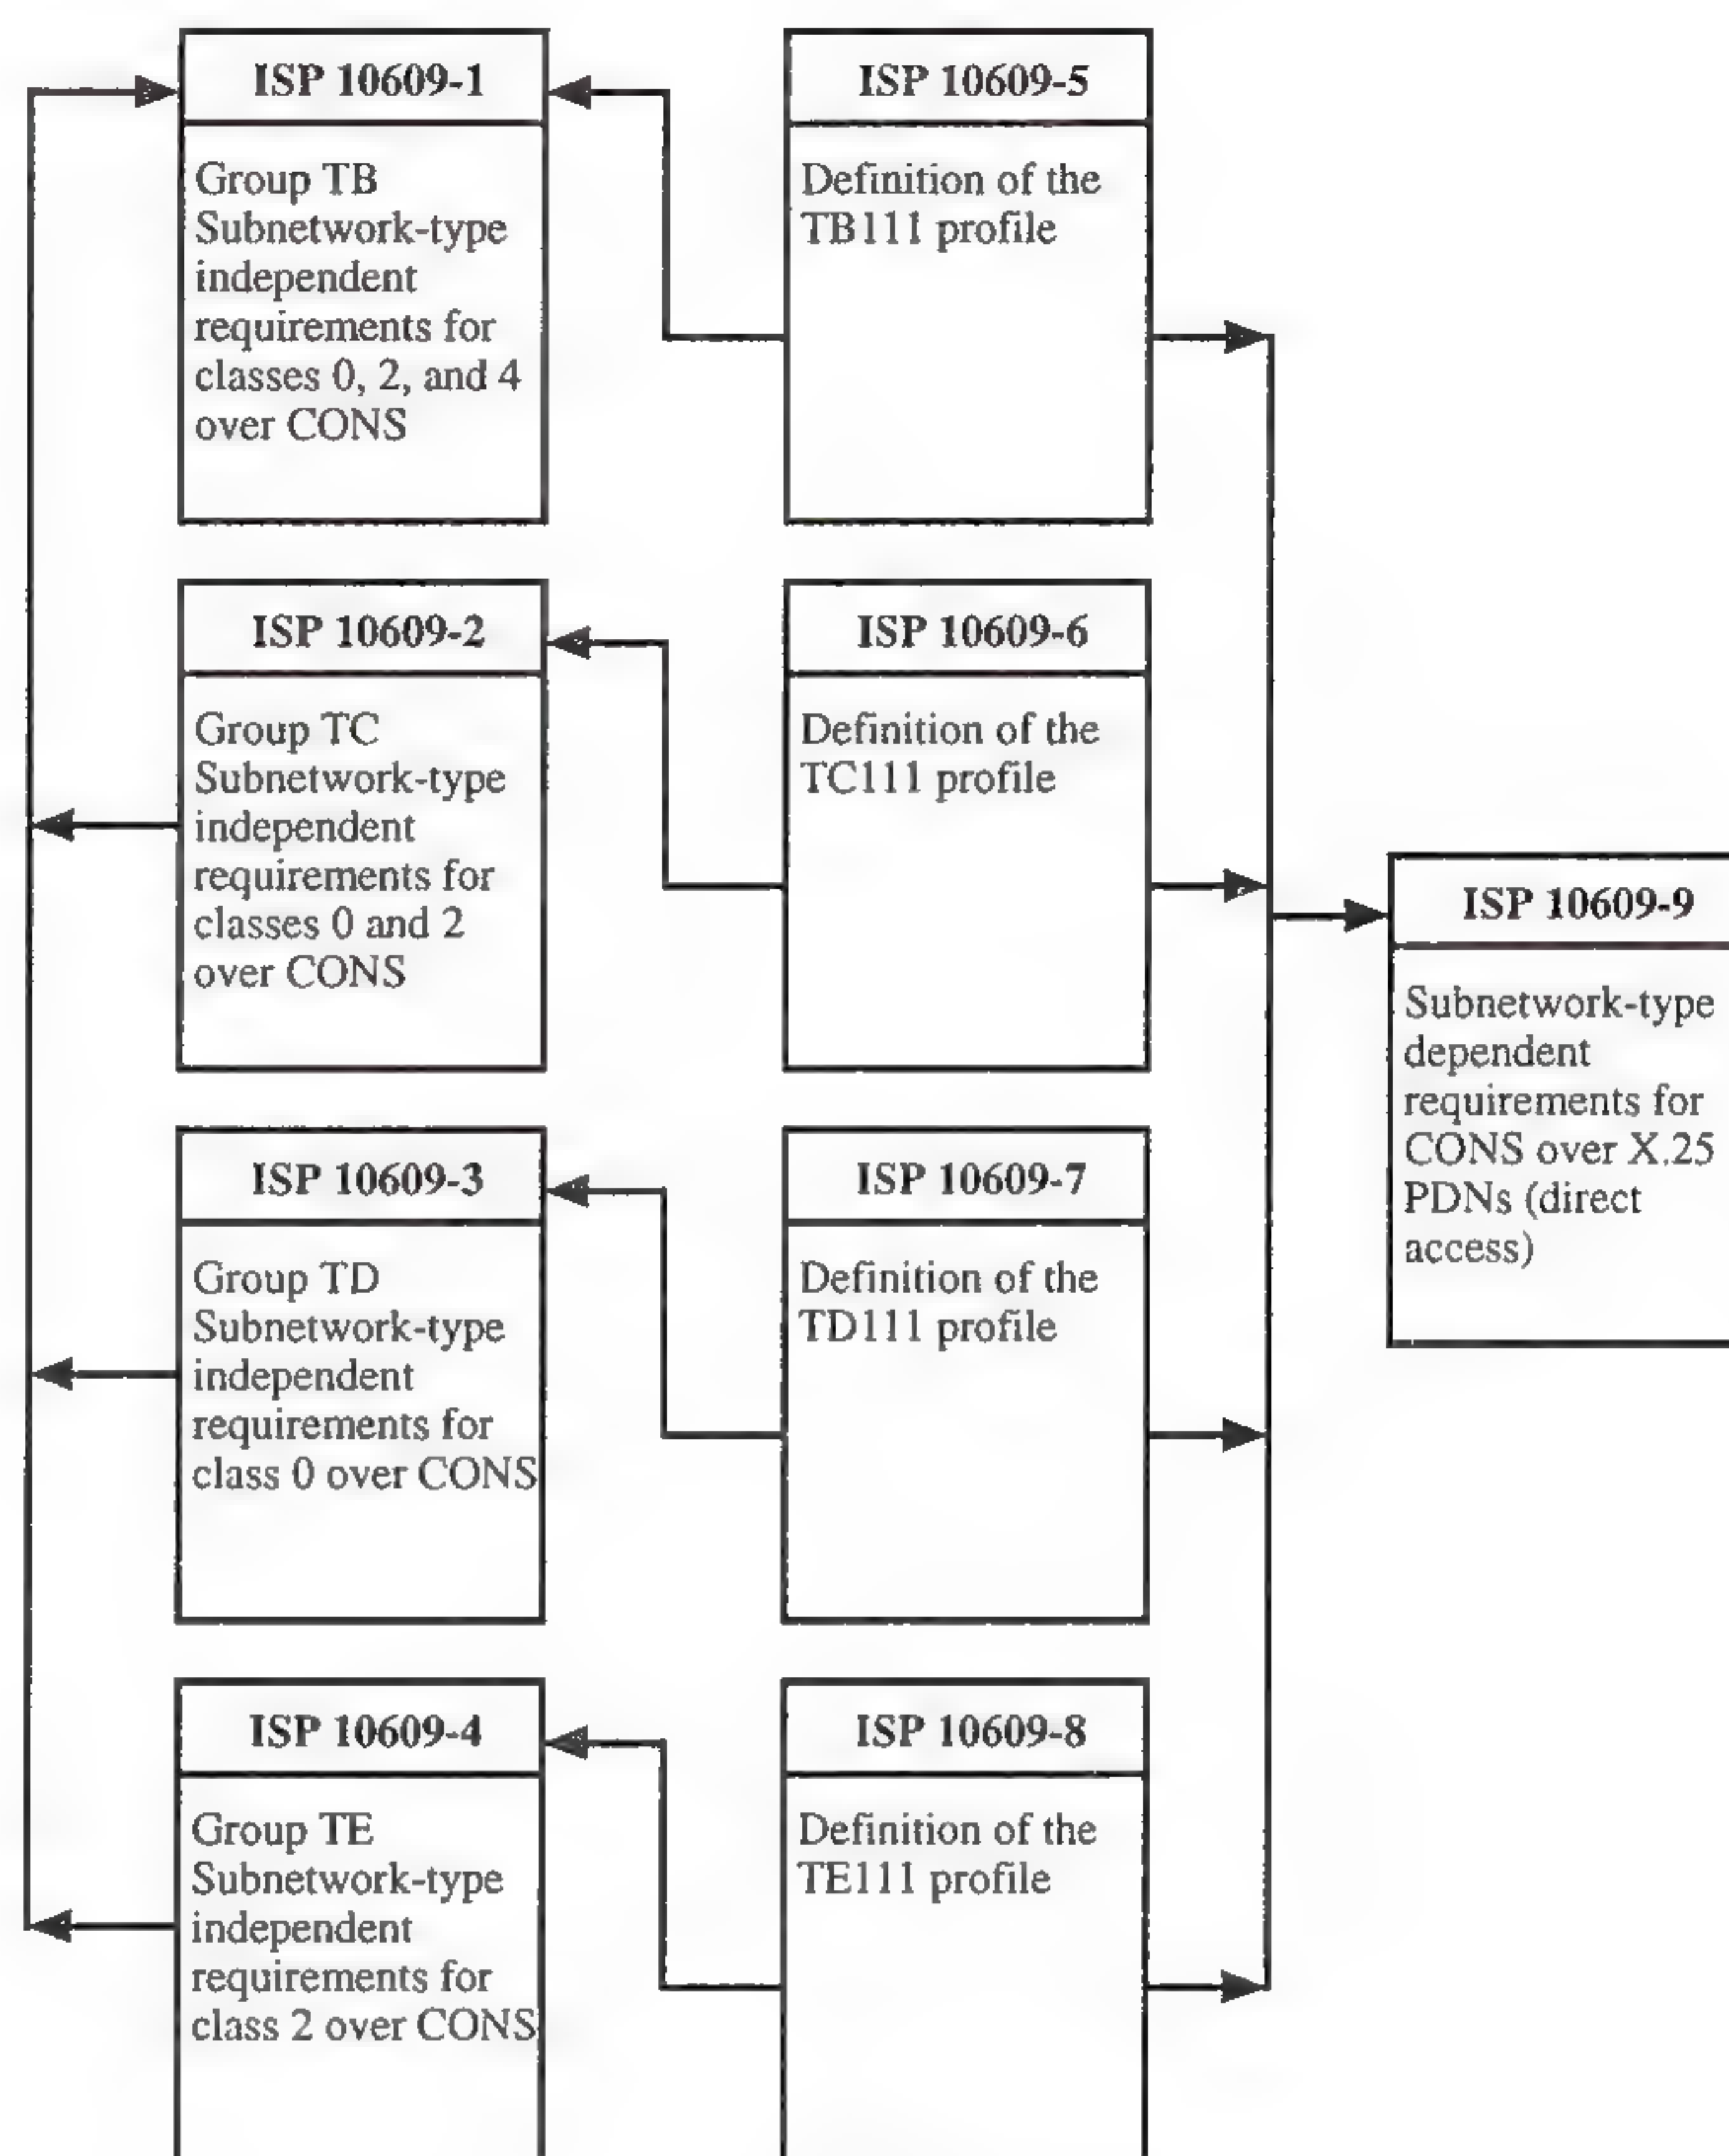


Figure 3: Multi-Part Structure of ISP 10609

Conformance

As with the implementation of a base standard, the implementation of a profile raises the issue of conformance. Indeed, with the development and standardization of the ISP approach, it is likely that most customers will primarily be concerned with the conformance of a profile implementation, since it is the profile that defines a set of protocols implemented to support a specific function. ISO addresses the issue of profile conformance in ISO 10000-1 and deals with profile conformance testing in CD 9646-6.

Profile conformance

A profile may contain conformance requirements that are more specific and limited in scope than those of the base standards to which it refers. For example, when a feature is associated with an allowed parameter value range, the profile can only adopt the same value range as that allowed by the base standard, or a subset of that range.

The key document used to define the conformance characteristics of a specific implementation is the *ISP Implementation Conformance Statement* (ISPICS). The ISPICS is in turn related to and derived from two sets of documents: the PICS proformas of the base standards and the *ISPICS Requirements List* (IPRL) that is part of the ISP.

The purpose of the IPRL is to provide a revised version of the conformance requirements of the constituent base standards of the profile.

International Standardized Profiles (continued)

The conformance requirements in the base standards (defined in the PICS proformas) relate to the conformance requirements in the profile (defined in the IPRL) in the following way:

- 1. Mandatory requirements in the base standards: Remain mandatory in the profile
- 2. Conditional requirements in the base standards: Remain conditional in the profile.
- 3. Optional requirements in the base standards: may be assigned to one of the following categories in the profile:
 - *Mandatory*: support may be made mandatory
 - *Optional*: support may remain optional
 - *Conditional*: optional requirements may be made conditional within the profile
 - *Out of scope*: optional requirements that are not relevant to the profile. For example, functional units of layer (n-1) that are unused by layer (n) in the context of the profile.
 - *Excluded*: The use of an optional feature may be prohibited in the context of the profile. This should only be used to restrict the dynamic behavior in terms of the transmission of protocol elements.
- 4. Non-applicable features in the base standards: These are features that are logically impossible according to the base standard. These remain non-applicable in the profile.
- 5. Excluded requirements in the base standards: Remain excluded in the profile.

Figure 4 illustrates this mapping. The only difference between static and dynamic conformance requirements is the ability to exclude optional dynamic requirements.

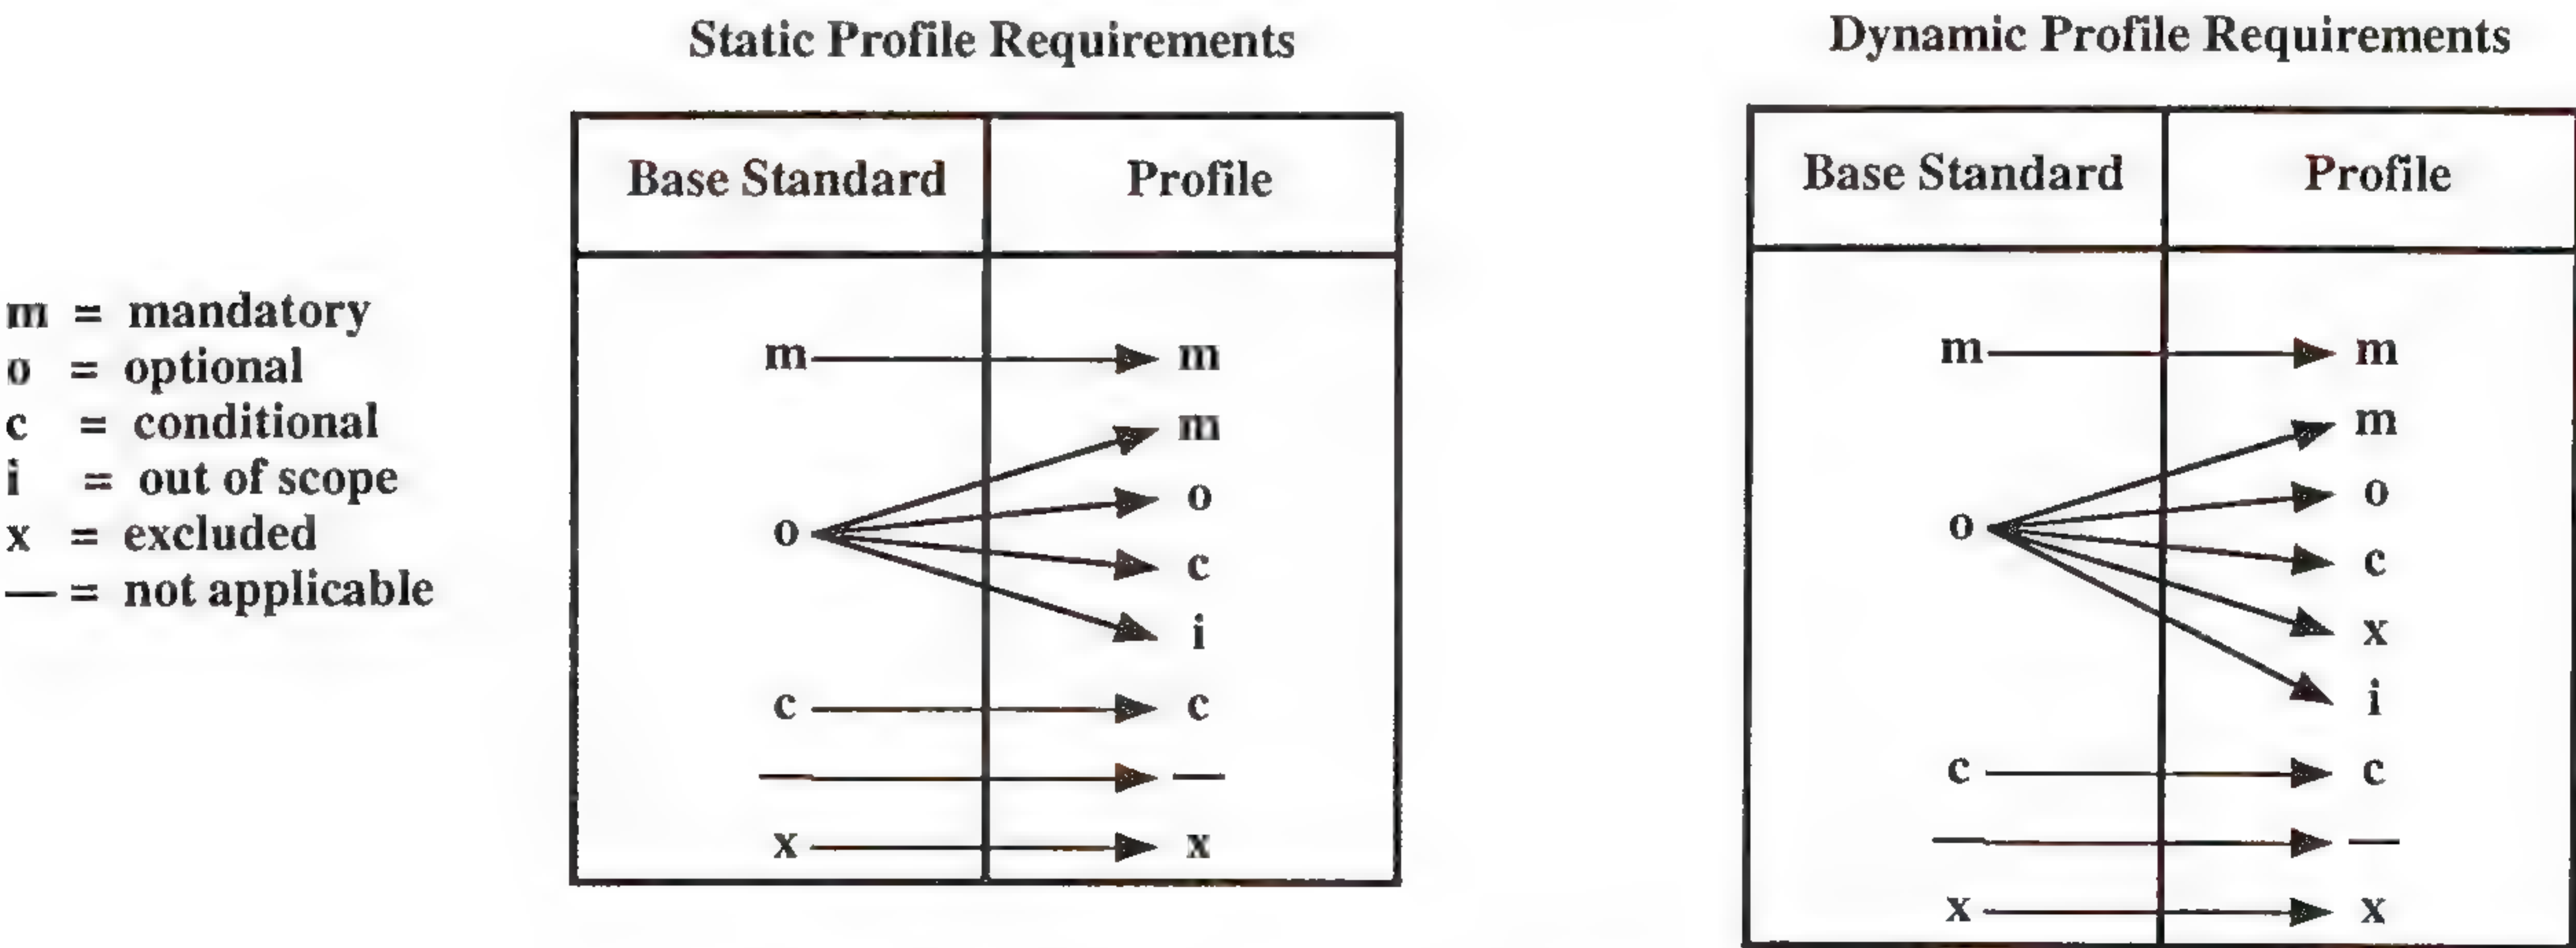


Figure 4: Static and Dynamic Profile Requirements

The relationship between and ISPICS and a PICS is as follows: A specific PICS proforma is associated with a specific base standard; it states the conformance requirements for the base standard in terms of mandatory, optional, conditional, and excluded features. The PICS proforma serves as a questionnaire to be filled out by the supplier of a specific implementation of the base standard.

The filled-out questionnaire is the PICS; it states the features supported by the implementation and the parameter ranges supported. A specific IPRL is associated with a profile; it states the conformance requirements for the profile in terms of mandatory, optional, conditional, and excluded features. An IPRL is provided for each profile in an ISP that contains multiple profiles. The IPRL specifies the profiles constraints on the supported features listed in the PICS proformas of the relevant base standards. The ISPICS consists of the set of PICS produced in accordance with the IPRL. For each base standard, the ISPICS states the features supported by the implementation and the parameter ranges supported.

Profile test specification

The development of conformance tests for a profile is based on the conformance test methodology for base standards. All test methods defined for testing base standards apply to profiles.

As with individual base standards, a preliminary step in testing a profile is to develop an abstract test suite. In the main, all that is required is to select the relevant test cases from the standardized abstract test suites for the base standards. A new test case is created for any profile-specific conformance requirement for which a testable test purpose is defined.

The defining document for profile testing is the *Profile Test Specification* (PTS). The PTS is applicable to only one profile; an ISP containing more than one profile will require a corresponding number of PTS documents.

A single profile test specification consists of three parts. The profile description and its profile requirements list forms the first part of the document. The PTS summary provides an overview of the test specification, with references to the base standards. Finally, the profile specific test specification is an optional part that holds additional test purposes and test cases beyond those defined for the base standards.

In the case of a ISP with multiple parts and multiple profiles, there is one part for each profile description/IPRL and one part for each PTS summary. Finally, there is a multi-profile test specifications part.

A PTS summary has the following format. Section 1 applies to the profile and the PTS as a whole, and contains general information relative to the profile. Sections numbered 2.N each apply to one protocol of the profile. It identifies the protocol and its PICS proforma, as well as the abstract *Test Suite Structure* (TSS) and *Test Purposes* (TP). Each abstract test suite for a protocol is referenced in a subsection, numbered 2.N.m. Section 3 provides information that is in addition to that related to that provided in the base standards; these are the additional test cases unique to the profile.

[This article is based on material in Bill Stallings' *Networking Standards: A Guide to OSI, ISDN, LAN, and MAN Standards*, © 1993 by Addison-Wesley. Used with permission. —Ed.]

WILLIAM STALLINGS is president of Comp-Comm Consulting of Brewster, MA and has twenty years of experience in data and computer communications. His clients have included major corporations and government agencies in the United States and Europe. Prior to forming his own consulting firm, he has been vice president of CSM Corp., a firm specializing in data processing and data communications for the health care industry. He has also been Director of systems analysis and design for CTEC, Inc., a firm specializing in command, control, and communications systems. He holds a PhD from M.I.T. in Computer Science and a B.S. from Notre Dame in electrical engineering. He is a frequent lecturer and the author of numerous papers and a dozen books on networking and computers. He can be reached at 72500.3562@compuserve.com.

The Message Send Protocol

by Billy Barron, University of North Texas

Introduction

For years, BITNET has had a method, known as TELL to VM users and SEND to VMS users, to send one line interactive messages to another user. Likewise, UNIX has the “write” command though it only works on the localhost. On the other hand, traditionally, TCP/IP has not had this functionality. Back in 1990, Russ Nelson wrote RFC 1159 for the experimental *Message Send Protocol* to add this functionality to TCP/IP.

Purpose

The Message Send Protocol can serve many different purposes. The first is that it is much simpler to implement with less system and network overhead than “talk.” In my experience, “talk” is inconvenient to use in some cases. The first is that I often have the need to send people one line messages and do not really need or expect a response from them. An example is a message to one of my operators saying “You forgot to start backups. Get going.” The other time is when I am in a single window environment and would like to execute UNIX or VMS commands while talking to someone. In addition, Some sites could use this protocol to display operator messages in a distributed environment. Finally, the protocol could be used for mailing list manipulation much like is currently done with BITNET LISTSERV.

RFC 1159: Version one

The initial protocol is very small and is documented in what has to be one of the shortest RFCs in existence. The Message Send Protocol of messages uses TCP or UDP port 18. The TCP and UDP versions both guarantee delivery in different ways though there is no method of preventing duplicate delivery. The packet format is:

A	Username	NULL	Terminal	NULL	Message	NULL
---	----------	------	----------	------	---------	------

The character “A” is the protocol version. The *Username* field is the name of the user the message is being sent to. If this field is empty, the message goes to the system console. This field and all of the remaining fields are null-terminated. The *Terminal* field is the name of the terminal to send the message to. If no terminal is specified, the RFC specifies that “the right terminal” is chosen. I guess this is the same terminal as chosen by “write.” The last field is the actual message being sent. The total length of the packet must be 512 octets or less.

RFC 1312: Version two

During the first part of 1992, the protocol was updated. The new packet format is:

B	RCPT Username	NULL	RCPT Terminal	NULL	Message	NULL
---	---------------	------	---------------	------	---------	------

Sender Name	NULL	Sender Terminal	NULL	Cookie	NULL	Sig	NULL
-------------	------	-----------------	------	--------	------	-----	------

The first four fields are almost the same as in RFC 1159. The protocol revision changed to the character “B”. The recipient terminal also now allows the value of “*” for messages to all terminals of the recipient username. The *Sender Name* is the username of the person who sent the message and the sender terminal is the name of the terminal the message was sent from. The *Cookie* is basically a message ID tag to detect and eliminate duplicate messages. Finally, the *Signature* field is for identifying the send for security reasons. The RFC does not define the method by which this field is used and says that a new RFC with the security feature will be issued at some point in the future.

Related work

Another change is in the method used for acknowledging received messages. A negative acknowledgement is now allowed. Positive acknowledgement messages can carry information about which user and terminal the message was actually delivered to.

SMTP, as defined in RFC 821, defines a method of sending a message to a user's terminal through port 25 via the SMTP SEND command. The implementation of this features is optional and most vendors do not implement it. I tested *IDA Sendmail*, *MMDF*, *Charon gateway*, *PMDF*, and *FAL* implementations of SMTP on my campus and found that none of them supported the SEND command. Also, as RFC 1312 points out, implementing the SMTP SEND command for sending interactive messages is a lot of work and is a high overhead protocol for the purpose.

In May 1988, Jim Frost defined the *User-to-User Message Transfer Protocol* (UMTP). UMTP is very similar to Version 1 of the Message Send Protocol, but there are a few differences. UMTP allows the ability to source route messages. The only use I see of this function is to traverse firewalls. The other major difference is that UMTP allows broadcasts to all the users on a node. UMTP does not have a regularly assigned TCP port number, which could lead to problems with widespread use of it.

Implementations

The SMTP SEND method has been implemented with a client program called *rsend*. The MSEND program by Jim Frost implements UMTP and has been previously posted to the alt.sources newsgroup. MSEND works with several different types of UNIX machines.

Version 1 of the Message Send Protocol has been implemented by Hebrew University as a program known as ISEND. ISEND works with many UNIX machines and VAX/VMS Systems running certain TCP/IP packages. ISEND is available for anonymous FTP on host VMS.HUJI.AC.IL in the local directory. No true implementations of Version 2 are publicly available as far as I know. The ISEND server has been modified to accept Version 2 packets, but it ignores the extra fields. ISEND should not be hard to modify to work with Version 2 of the protocol specification. The question is who is going to do the work.

The future

With the possible death of BITNET in the next few years, users who are migrating from BITNET to the Internet will want the same functionality as they have had on BITNET. To make them happy, a Message Send Protocol server needs to be tied into the UNIX version of LISTSERV. If this is not done, support personnel at BITNET sites will have to do significantly more retraining of users.

Conclusion

An interactive single message protocol is badly needed on the Internet. The Message Send Protocol, though experimental, is a good mechanism by which to do this.

References

- [1] Frost, Jim, "MSEND—Immediate message sending program for UNIX machines," 1991.
- [2] Frost, Jim, "User-to-User Message Transfer Protocol (UMTP)," 1988.
- [3] Nelson, Russ, "Message Send Protocol," RFC 1159, 1990.
- [4] Nelson, Russ & Geoff Arnold, "Message Send Protocol 2," RFC 1312, 1992.
- [5] Postel, Jon, "Simple Mail Transfer Protocol," RFC 821, 1982.
- [6] Thomas, Eric. "Revised List Processor (LISTSERV@FRECP11), Release 1.5d," École Centrale de Paris, 1986.

BILLY BARRON received his B.S. (1986) from North Texas State University and M.S. (1988) from University of North Texas. For the past five years he has worked for the University of North Texas in various capacities and now is the VAX/UNIX Systems Manager running VAX and various UNIX computer systems.

Inside an IETF Plenary Meeting

by Gary Scott Malkin, Xylogics

Introduction

The most recent *Internet Engineering Task Force* (IETF) Plenary meeting was held in Washington D.C. from November 16–20, 1992. This was the Silver Anniversary Plenary meeting. Statistics tell us that 38% of those people who attend the Plenary meeting do so for the first time. Since the attendance at these meetings has grown so much over the last few years, it seems time to describe how they operate. This way, a newcomer to the IETF will not be quite so overwhelmed and will be able to get much more out of the meeting.

Background

Here's the official definition of the IETF:

The IETF is the protocol engineering, development, and standardization arm of the *Internet Architecture Board* (IAB). Its mission includes:

- Identifying, and proposing solutions to, pressing operational and technical problems in the Internet;
- Specifying the development or usage of protocols and the near-term architecture to solve such technical problems for the Internet;
- Making recommendations to the IAB regarding standardization of protocols and protocol usage in the Internet;
- Facilitating technology transfer from the IRTF to the wider Internet community; and
- Providing a forum for the exchange of information within the Internet community between vendors, users, researchers, agency contractors, and network managers.

In order to accomplish this mission, the IETF holds three Plenary meetings a year. In addition, the Working Groups within the IETF may meet between Plenaries, in person or by audio/video conference.

IETF structure

The IETF has a structure which has evolved over the years. Within the IETF there are nine functional *Areas*: Applications, Internet Services, Network Management, Operational Requirements, OSI Integration, Routing, Security, Transport and Services, and User Services. Each Area has at least one *Area Director*. There is also an Area Director responsible for Standards Management. The Area Directors, along with the IETF Chair, make up the *Internet Engineering Steering Group* (IESG). The IESG provides the first technical review of Internet Standards. It is also responsible for the day-to-day "management" of the IETF.

Each Area has several *Working Groups*. A Working Group is a group of people who work under a charter to achieve a certain goal. That goal may be the creation of an informational document, the creation of a protocol standard, or the resolution of problems in the Internet. Most Working Groups have a finite lifetime. That is, once a Working Group has achieved its goal, it disbands.

Areas may also have *Birds of a Feather* (BOF) groups. They generally have the same goals as Working Groups, except that they have no charter and usually only meet once or twice. BOFs are often held to determine if there is enough interest to form a Working Group.

Membership and mailing lists

There is no official membership in the IETF. The closest thing there is to membership is being on the IETF *mailing lists*. There are two main IETF mailing lists, one for announcements and one for discussions of cosmic significance. To join the announcement list, send an e-mail message to: `ietf-announce-request@cnri.reston.va.us`.

To join the discussion list, `ietf@cnri.reston.va.us`, send a message to `ietf-request@cnri.reston.va.us`. Do not, ever, for any reason, send requests to join or leave a list to the list itself. Always use the “-request” address.

There is also no formal membership for the Working Groups within the IETF. Each Working Group has an e-mail discussion list and a “-request” address for administrivia. To be a Working Group “member” join the list and, if possible, attend the Working Group’s meetings.

Registration

The *IETF Secretariat* (the people who make the Plenary meetings work) send Plenary meeting announcements to the IETF announcement list beginning about three months before the meeting. Within the IETF meeting announcement is a Registration Form and complete instructions for registering, including, of course, the cost. The Secretariat highly recommends that attendees preregister. Early registration, which ends about one month before the meeting, carries a lower registration fee. As the size of the meetings has grown, so has the length of the lines at the registration desk. Fortunately, there are three lines: the “preregistered and prepaid” line (which moves very quickly); the “preregistered and on-site payment” line (which moves a little more slowly); and the “registration and on-site payment” line (take a guess).

Registration is open all week. However, the Secretariat highly recommends that attendees arrive for early registration, beginning at 6:00 pm. (meeting local time), on the Sunday before the meeting. Not only will there be fewer people, but there will also be a reception at which people can get a byte (!) to eat. If the registration lines are long, one can eat first and try again when the lines are shorter. Newcomers are encouraged to attend the IETF Orientation, Sunday at 4:30 pm.

Registered attendees (and there isn’t any other kind) receive a Registration Packet. It contains a general orientation sheet, the At-A-Glance sheet, a list of Working Group acronyms, the most recent Agenda, and a name tag. The At-A-Glance is a very important reference and is used throughout the week. It contains Working Group/BOF room assignments and a map of room locations. Attendees who prepaid will also find their receipt in their packet.

What to expect

One thing which the Plenary meeting announcement doesn’t address is the IETF dress code. Since attendees must wear their name tags, they must also wear shirts. Pants are also highly recommended. Most people wear T-shirts, jeans and sandals. A very few wear suits. The general rule is “dress for the weather.”

Some of the people at the IETF will have a little colored dot on their name tags. A few people have more than one. These dots identify people who are silly enough to volunteer to do a lot of extra work. The colors have the following meanings:

- *Red*: IAB member
- *Yellow*: IESG member
- *Blue*: Working Group/BOF chair
- *Green*: Local host

Inside an IETF Plenary Meeting (*continued*)

Local hosts are the people who can answer questions about the terminal room, and restaurants and points of interest in the area.

Internet Access

For those who go into withdrawal if they can't read their e-mail, the local hosts provide a terminal room with Internet access. In general, the connectivity is excellent. This is entirely due to the Olympian efforts of the local hosts, and their ability to beg, borrow and steal. The people and companies who donate their equipment, services, and time are to be heartily congratulated and thanked.

Finally there is the Social Event, which has become something of a tradition at the Plenary meetings. The local hosts handle all of the arrangements and send announcements to the IETF announcement list. Some Social Events are high-tech, like a tour of the Stanford Linear Accelerator. Others include dinner cruises and trips to museums, pipe organs and art galleries. Newcomers are encouraged to attend the Social Event. Everyone is encouraged to wear their name tags. The social is designed to give people a chance to meet on a social, rather than a technical, level.

Important e-mail addresses

There are some important IETF e-mail addresses with which everyone should be familiar. They are all located on Internet host `cnri.reston.va.us` (e.g., `ietf-info@cnri.reston.va.us`).

- `ietf-info`: *General queries about the IETF*
Greg Vaudreuil, Megan Davies and Cynthia Clark
- `ietf-rsvp`: *Queries about meeting locations and fees, e-mailed Registration Forms*
Debra Legare
- `proceedings`: *Queries about previous Proceedings availability, orders for copies of the Proceedings*
Debra Legare
- `ietf-announce-request`:
Requests to join/leave IETF announcement list
Cynthia Clark
- `ietf-request`:
Requests to join/leave IETF discussion list
Cynthia Clark
- `internet-drafts`:
Internet-Draft submissions
Cynthia Clark
- `iesg-secretary`:
Greg Vaudreuil

IETF proceedings

The IETF Proceedings are compiled in the two months following each IETF meeting. The Proceedings usually start with a message from Phill Gross, the Chair of the IETF. Each contains the final (hindsight) Agenda, an IETF overview, a report from the IESG, Area and Working Group reports, network status briefings, slides from the protocol and technical presentations, and the attendees list. The attendees list includes an attendee's name, affiliation, work phone number, work fax number, and e-mail address, as provided on the Registration Form.

A copy of the Proceedings will be sent to everyone who registered for the IETF. The cost is included in the registration fee. The Proceedings are sent to the mailing addresses provided on the Registration Forms.

For those who could not attend a Plenary, but would like a copy of the Proceedings, send a check for \$35 (made payable to CNRI) to:

Corporation for National Research Initiatives
Attn: Accounting Department—IETF Proceedings
1895 Preston White Drive,
Suite 100
Reston, VA 22091

Please indicate which Plenary Proceedings you would like to receive by specifying the meeting date (e.g., July 1992) or meeting number and location (e.g., 24th Plenary in Boston). Availability of previous Plenary Proceedings is limited, so check *before* sending payment.

Be prepared

This topic was saved for last, but it is by far the most important. As the IETF grows, it becomes more and more important for attendees to arrive prepared for the Working Groups meetings they plan to attend. This doesn't apply only to newcomers; everybody should be prepared to make the most out of the week.

Being prepared means having read the documents which the Working Group or BOF Chair has distributed. It means having followed the discussions on the Working Group's mailing list or having reviewed the archives. For the Working Group/BOF Chairs, it means getting all of the documents out early enough (i.e., several weeks) to give everybody time to read them. It also means announcing an agenda and sticking with it.

At the Chair's discretion, some time may be devoted to bringing new Working Group attendees up to speed. In fact, long-lived Working Groups have occasionally held entire sessions which were introductory in nature. As a rule, however, a Working Group is *not* the place to go for training. Observers are always welcome, but they must realize that the work effort cannot be delayed for education. Anyone wishing to attend a Working Group for the first time might seek out the Chair prior to the meeting and ask for some introduction.

Another thing, for everybody, to consider is that Working Groups go through phases. In the initial phase (say, the first two meetings), all ideas are welcome. The idea is to gather all the possible solutions together for consideration. In the development phase, a solution is chosen and developed. Trying to reopen issues which were decided more than a couple of meetings back is considered bad form. The final phase (the last two meetings) is where the "spit and polish" are applied to the architected solution. This is not the time to suggest architectural changes or open design issues already resolved. It's a bad idea to wait until the last minute to speak out if a problem is discovered. This is especially true for people whose excuse is that they hadn't read the documents until the day before a comments period ended.

Time at the IETF Plenary meetings is a precious thing. Working Groups are encouraged to meet between IETF meetings, either in person or by video or telephone conference. Doing as much work as possible over the mailing lists would also reduce the amount of work which must be done at the meeting.

Advice from an Old Timer

I've attended 12 of the last 14 IETF meetings. I remember that, for the first few, I was hesitant to introduce myself to people whose names I knew from the tops of RFCs. It took me a while to realize that they were Internet People just like me (albeit with a lot more experience).

Inside an IETF Plenary Meeting (*continued*)

IETFers are very friendly people. They're always willing to stop and chat (unless they're on their way to get something to eat). You should never be afraid to walk up to someone and introduce yourself. A tremendous amount of work is accomplished between Working Group sessions by people who gather in the halls. Many lunches, dinners and late night get-togethers devolve into technical discussions (to the consternation of some).

The important thing to remember is that the IETF exists to promote the growth of the Internet. The primary reason the Internet, and the protocols which run on it, are so successful is the openness of the process and the belief in "Rough Consensus and Working Code." If you'd like to contribute, or just listen in, you'll always be welcome to do so.

References

- [1] Gross, P., "The Internet Engineering Task Force (IETF)," *ConneXions*, Volume 2, No. 10, October 1988.
- [2] Postel, J., "An overview of the Internet Activities Board," *ConneXions*, Volume 1, No. 8, December 1987.
- [3] Cerf, V., "Internet Activities Board," RFC 1160, May 1990.
- [4] Casner, S., & Deering, S., "First IETF Internet Audiocast," *ConneXions*, Volume 6, No. 6, June 1992.
- [5] Lottor, M., "Internet Growth (1981–1991)," RFC 1296, January, 1992.
- [6] Solensky, F., "The Growing Internet," *ConneXions*, Volume 6, No. 5, May 1992, pp 46–48.
- [7] Marine, A., "How Did We Get 727,000 hosts?" *ConneXions*, Volume 6, No. 5, May 1992, pp 49–51.
- [8] Crocker, D., "The ROAD to a New IP," *ConneXions*, Volume 6, No. 11, November 1992.
- [9] Chapin, L. (ed.), "The Internet Standards Process," RFC 1310, March 1992.

GARY SCOTT MALKIN received his B.A. in Computer Science from Boston University in 1983 and his M.S. in 1992. After almost five years working for Spartacus (he was the model for Fibronics' two advertisements), and almost four years elsewhere, he is now happy as a Principal Software Engineer at Xylogics in Burlington Mass. Gary is an active member of the Internet Engineering Task Force. He is a member of the User Services Advisory Council and has chaired several Working Groups. Currently, he is the chair of the RIP-2 Working Group and co-chair of the Internet User Glossary Working Group. He is an author of several RFCs, including the "Questions and Answers" FYI RFCs. E-mail: gmarkin@xylogics.com.

On Paper

We are pleased to announce that, starting with this issue, *ConneXions* will be printed on the recycled version of the paper we've been using for several years (Simpson, Gray Mustang, 60#). While recycled paper is actually slightly more expensive than its ordinary counterpart, we're happy to do our small part for the environment.



Printed on recycled paper

NREN Politics Thicken

by Mike Roberts, EDUCOM

Introduction

The wild-eyed success of INTEROP 92 Fall at Moscone Center in San Francisco, reported previously in this publication, underlines the extent to which double and triple digit growth rates are outrunning the capacity of Internet political, economic and technical institutions to scale up to meet the challenge of planning and implementing a very large worldwide network.

Three separate points of view on the future of the Internet are beginning to emerge. Pursued literally, the different approaches could lead to very different Internet outcomes.

Carriers

Communications carriers, national and international, are coming alive to the business opportunities represented by an Internet connecting fifty to one hundred million computers by the end of the decade. Having had little to do with the Internet until recently, except for the furnishing of private lines, their interest is not matched by agreed upon strategies. But their voices are similar on two points, don't regulate it and don't let government, or other public sector organizations, have very much to do with it.

Public sector

The traditional research and higher education constituencies of the Internet find themselves suddenly surrounded by many new non-profit friends, including libraries, primary and secondary schools, state and local government agencies, and others who want to "put the *E* in NREN." This is an aggregation of strange bedfellows whose historical strength is in diversity, not collective action. In addition to their troubles in forging an effective coalition for the national network from among their own number, they fear continuing domination by federal research agencies on the one hand, and a sudden swing to full scale commercialization on the other hand.

Agencies

The "big five" research agencies (DARPA, Energy, HHS, NASA and NSF), who take, and even deserve, a lot of credit for getting us to where we are, see their baby being carried off by ambitious politicians intent on the rhetoric of the "information highway." The heads-down "we're only doing research" approach isn't working anymore, but the federal establishment has no more of a new game plan than any of the other players.

Simply put, the carriers want us to trust them because they know how to "do" communications; the public sector wants its values respected and preserved; and what the feds want depends on who you talk to.

What next?

A disinterested observer might suggest that all these folks have more to gain by working together than by fighting. But after a summer marked by community rejection of the NSFNET draft solicitation, a barroom brawl among competing IETF factions, and nearly everyone coming up short in their budget battles, tempers are fragile and skins are thin.

The challenge to Bill Clinton and Al Gore is to define a national network agenda that transcends day to day difficulties and puts us all in harness and pointed in the right direction.

MIKE ROBERTS is Vice President for Networking at EDUCOM, a 600 member association of colleges and universities with common interests in information technology. The EDUCOM Networking and Telecommunications Task Force, a group of sixty universities and corporations, of which he is the staff director, has been active in planning and advocacy for the NREN. E-mail: roberts@educom.edu

SLIP at the University of Minnesota

by Craig A. Finseth, University of Minnesota

Need Our users have been requesting dialup SLIP (*Serial Line IP*) access to our network for some time. Three departments of the University, CIS/Networking Services, Telecommunications Services and Distributed Service and Planning, have jointly set up this service.

Dialup SLIP service allows a user to dial into the University network and act as a full network device. While they can use Telnet to log into hosts, they can also:

- Directly FTP to and from other computers,
- Can use *PopMail* to read their mail,
- Can read News, and
- Can access other protocols.

Users may connect to hosts both within the University and anywhere on the Internet.

Background While designing this system in the fall of 1991, I checked with a wide range of people around the Internet, asking for examples of what other places do as far as SLIP service. No one had any ideas. Hence, the system and its security/authentication requirements were developed more-or-less in a vacuum.

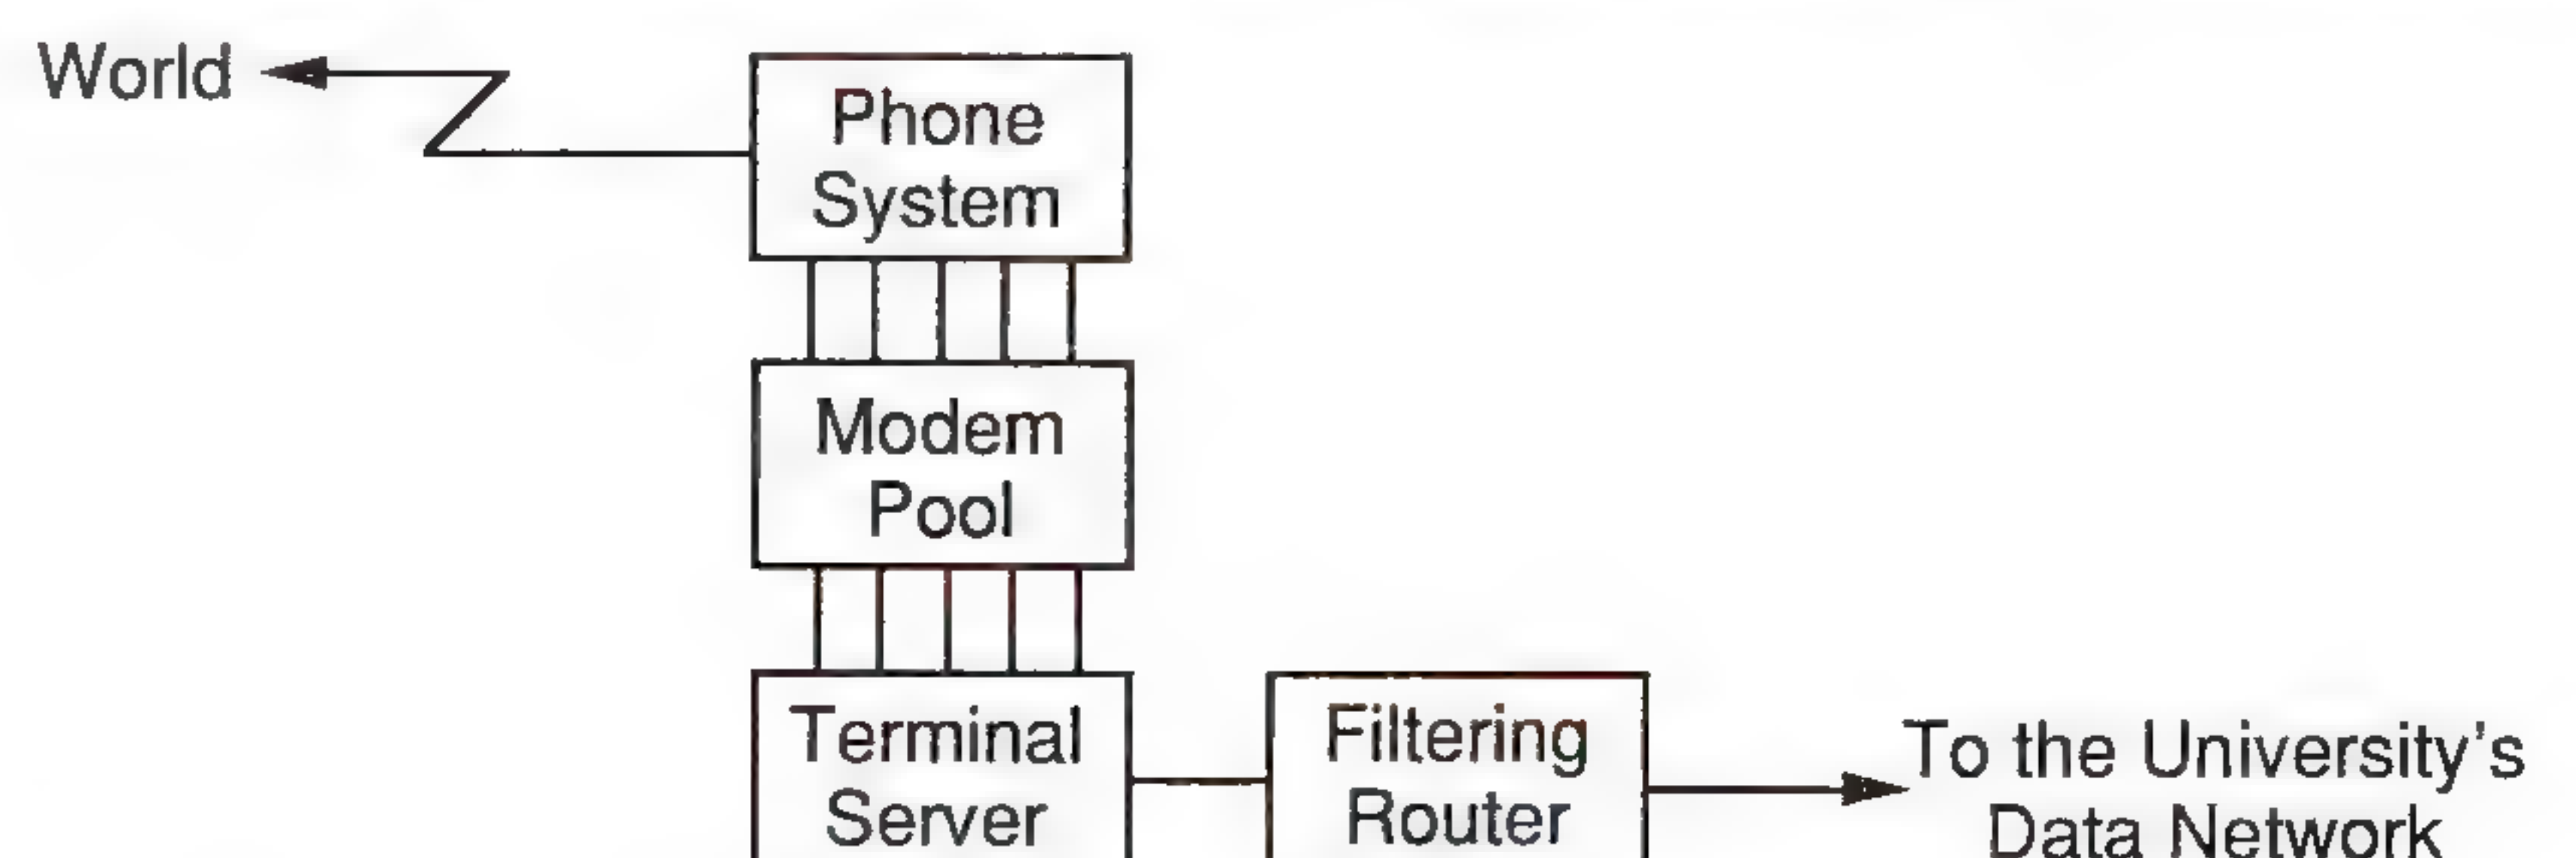
Basic operation The user obtains SLIP client software for their Macintosh or IBM PC-compatible computer from the Distributed Services and Planning group. This software comes pre-configured with phone numbers, etc. It also uses BOOTP to automatically assign an IP address to the machine. A UNIX version of the software is also planned.

The user installs and runs this software on his or her system. The user then dials and connects to the Cisco terminal server. The terminal server asks that the user enter his or her e-mail address and a password. (Authentication is described fully in a later section.) Once given, the user is free to invoke SLIP and operate as a network device.

The required hardware is a Macintosh or IBM PC-compatible computer and 2400 or 9600 bps modem. Other computers can be used, but users must supply their SLIP client software.

Users are permitted to operate as clients. But, the access is via a modem pool, with each modem having a different IP address. Therefore, users may not operate servers using this service, as there is no assurance that subsequent calls will receive the same IP address.

Hardware Users dial into a modem pool on the campus telephone system. The modem pool is connected to a Cisco terminal server, which is in turn connected to the network through a filtering router:



Our phone system uses rotaries to find the next available modem. The modems will automatically negotiate speeds and compression, etc.

Authentication and security

All lines on this terminal server are set to 8-bit clean by default. The terminal server is intended for a variety of users (Emacs, Xmodem, etc.) besides SLIP traffic. As all use of the terminal server is authenticated, there are no outgoing routing restrictions. Currently there is no check for loop use of SLIP (i.e., using the SLIP access to get another SLIP access to hide origination of the connection), but the authentication step substantially reduces the risk of that type of attack.

The filtering router is a Cisco IGS/L. It has a number of network-layer filters that prevent SLIP hosts from interfering with normal network operation. In general, all routing protocols and most harmful ICMP responses are filtered out. This router is required because Cisco only supports filters on *outgoing* traffic, and the next hop is a large Cisco AGS+ with about 30 interfaces. It makes little sense to burden that router with the filters required for only one, slow speed channel.

Ethernet networks are used to interconnect the terminal server, filtering router, and the data network.

We use the TACACS protocol supported by Cisco for authentication. This is a UDP-based protocol originally used by the ARPANET TACs (*Terminal Access Controllers*). Users must enter their e-mail address and a password. The terminal server forwards this information (using TACACS) to an authentication host.

This host (a Sun) logs the request and extracts the host portion of the e-mail address. It looks this portion up in a table of approved hosts. If not present in the table, the authentication request is rejected. If the host is present in the table, the user name portion and password are forwarded to that host for verification. The target host responds with a yes or no. The authentication host logs the outcome (no matter what it is), and returns a "yes" or "no" to the terminal server.

A request can be rejected for a variety of reasons:

- Misformatted or invalid request.
- Host not on the authorized list.
- User name on a special "hit list."
- Zero length password.
- Missing or negative response from the host.

This scheme allows us to have what amounts to a "unique identifier," but without having to create a central list of accounts and passwords.

In order to have a host added to the authorized list, the system manager of that host must certify that there are no "guest" type accounts and that the manager will work with us to track down any security or verification problems.

We are very conscious of possible security problems. The terminal server and related equipment are physically in the access-controlled Telecommunication center. The authentication host is in the access-controlled CIS/Networking Services area. These areas are connected to each other and to the rest of the University with fibre optic links. Therefore, there are essentially no practical ways to monitor all TACACS traffic.

SLIP at the University of Minnesota (*continued*)

User TACACS traffic is exposed when it reaches the building that contains the authorized host. (But only TACACS requests for hosts in that building will be sent to that building.) Still, this exposure also exists for Telnet, FTP, and other protocols that send passwords as clear text, so the problem is not made any worse.

The TACACS servers on the authorized hosts have been modified to only respond to requests originating from the authentication host. Hence, a third party cannot use TACACS to “crack” accounts.

Each request and its disposition is logged by the authentication host. A program, run daily, analyzes these logs looking for patterns of excessive validation or excessive failed attempts. We can thus often take action to prevent problems with accounts *before* they have been broken into. It also logs the length of each connection, and if someone is tying up a line for an excessive amount of time, we can contact them.

Summary

We provide SLIP access to the network in a way that greatly benefits our users, and that we believe complies with Internet policies, prevents disruption of the network, and discourages attempts at “cracking.”

References

- [1] R. Coop, “SLIP Interoperability” *ConneXions*, Volume 6, No. 6, June 1992.
- [2] C. Partridge, “Dialup IP,” *ConneXions*, Volume 3, No 11, November 1989.
- [3] J. Romkey, “A Nonstandard for Transmission of IP Datagrams over Serial Lines: SLIP,” RFC 1005, June 1988.
- [4] J. Romkey, “SLIP: Serial Line IP,” *ConneXions*, Volume 2, No. 5, May 1988.
- [5] R. Hobby, “The Point to Point Protocol (PPP)—A new proposed standard Serial Line Protocol,” *ConneXions*, Volume 4, No. 4, April 1990.
- [6] V. Jacobson, “Compressing TCP/IP Headers for Low Speed Serial Links,” RFC 1144, February 1990.

CRAIG FINSETH hold B.S. degrees in Computer Science and Engineering and Philosophy from M.I.T. He is the author of *The Craft of Text Editing*. Since 1987, he has worked for the Minnesota Supercomputer Center and is a “founding” employee of the University of Minnesota’s Networking Services department. He has been involved with many network management and computer security projects, including the development of a state-of-the-art network management system designed to meet the needs of large networks. He has been involved with all aspects of designing, operating, and managing the University’s 10,000+ node network. He started and is currently overseeing the final stages of a two-year reconstruction of the University’s backbone network. His e-mail address is: `fin@unet.umn.edu`.

Call for Papers

The 18th Annual *Conference on Local Computer Networks* (LCN)—“The Conference on Practical Leading Edge Computer Networking”—will be held September 19–22, 1993 in Minneapolis, Minnesota, USA. The event is sponsored by the IEEE Computer Society, Technical Committee for Computer Communications. The program consists of Tutorials, Technical Paper Sessions, and Panel Discussions.

Theme The emphasis on this year's conference is on practical experience using local computer networks. This unique approach simulates a workshop environment and allows for an effective interchange among users, researchers, and vendors. Some of the primary goals of the conference are to enable those involved in the local computer network field to share experiences, lessons learned, and prototype data and analysis. Because of these objectives, papers based on experience are especially solicited. The focus of the 18th LCN will be Multimedia and Video Communications. Papers that cover these areas are explicitly sought and will be given preference.

Topics Sessions are being organized on:

- Multimedia on the LCN
- Real-Time Video Applications
- ATM
- High Speed Networks
- Bandwidth Allocation
- Compression Techniques
- Congestion Control & Avoidance
- Metropolitan Area Networks
- Internetworking
- Network Management
- Emerging Technologies
- Security
- Reliability, Availability, & Maintainability (RAM)
- Multimedia Applications
- HIPPI
- Fiber Channel Networking
- Gigabit Networks
- Isochronous Protocols
- Error Control Techniques
- High-Performance Protocols
- LAN/MAN/WAN Integration
- Standards
- Remote Monitoring
- FDDI and FDDI-II

Submissions All authors must submit 5 full copies of the full technical paper by mail or delivery service. *Do not submit complete papers via fax.* The first page must contain: title of the paper, author's names including affiliations, complete mailing address, telephone and Fax numbers, Internet or BITNET address, and a 250 word (maximum) abstract (double-spaced) in English to the Program Chair, at this address:

Dr. Kenneth Ocheltree, Program Chair
 IBM T.J. Watson Research Center
 Mail Stop H4-A24
 30 Saw Mill River Road
 Hawthorne, NY 10532
 Phone: +1 914-784-7903
 Fax: +1 914-784-6201
 E-mail: keno@watson.ibm.com

Important dates

Deadline for submission:	April 5, 1993
Acceptance:	June 28, 1993
Camera-ready copy due:	July 30, 1993

Novell Licenses NetWorld to Interop Company

NetWorld+INTEROP 94

Mountain View, California, December 18, 1992—Interop Company today announced that in connection with Novell, Inc. it will form a new event beginning in 1994 called *NetWorld+INTEROP 94*. The move is intended to benefit both attendees and exhibitors by producing a single event with an integrated exhibition and educational program designed to meet the industry's growing need to understand interoperability issues at all levels of computing and communication.

Under the terms of an agreement with Interop Company, Novell, Inc. of Provo, Utah, has licensed worldwide rights to use the name "NetWorld" and to manage the show beginning January 1994.

All levels of networked interoperability

"Business computing is becoming networked computing," said Dan Lynch, president and founder of Interop Company. "The issues buyers face are less centered on computing as such and more on how to make computers at all levels communicate and work together seamlessly—from desktop to data center. NetWorld® has an excellent reputation for serving those involved in network computing and INTEROP® has a proven track record for addressing enterprise computing customers. NetWorld+INTEROP 94 will address the industry's growing demand for educational, technically-relevant forums that reflect the latest developments—at *all* levels of networked interoperability."

"As enterprise computing and network computing become more integrated into global networking issues, it makes sense to provide an event where customers can learn what they need to know at one time and in one place," said Terri Holbrooke, vice president of corporate communications for Novell.

INTEROP

Interop Company has become known for its conference and exhibition, INTEROP, which is held twice a year domestically and will expand into Europe with a show in Paris in October 1993.

Unlike traditional trade show companies, which focus primarily on exhibits, Interop Company has been a leader in staging multi-faceted events where education and technology demonstrations are as important as exhibit floor-based marketing activities.

INTEROP 92 Fall, Interop Company's most recent event, was held in San Francisco, California, and attended by more than 55,000 people. In addition to the three-day exhibition where exhibitors modeled real-world interoperability by demonstrating their products on the INTEROPnet® network, the San Francisco event offered a rich educational program. There were more than 40 technology-based tutorials lead by the industry's leading academic and research professionals.

INTEROP 92 Fall also featured four "conferences within a conference." These sub-conferences were: Desktop INTEROP, focusing on LANs and PC networking; Global INTEROP, addressing issues of how to construct internet infrastructures; SNA INTEROP, for those coming from a mainframe perspective; and Executive INTEROP, to give corporate management attendees a macro view of interoperability.

Worldwide strategy

The NetWorld+INTEROP 94 concept reflects Interop Company's evolving worldwide strategy for the 1990's: to provide the computing and communications communities with a forum for understanding and demonstrating the technologies, products and interoperability issues that affect their businesses.

"Rather than react to issues in the industry, Interop Company has always been committed to anticipating them," said Lynch. "The growing popularity of INTEROP today is evidence that we've been on the right track with interoperability. NetWorld+INTEROP 94 is part of our strategy to anticipate the future."

Since December 1990, Interop Company has been backed by its parent company, Ziff-Davis, known for its visionary and innovative approach to helping its audiences keep up with the latest developments in the Information Age.

"We're happy to see Interop Company expand its ability to serve networking customers with events such as NetWorld+INTEROP 94," said William Lohse, president of Ziff-Davis' newly formed Ziff-Davis Exhibitions and Conference Company. "Interop is modeling the trade show of the 21st century and has already proven its approach. We're looking forward to helping many more customers with this new show."

About Interop Company

Interop Company is an educational services company dedicated to bringing the latest ideas from researchers, analysts and vendors to the user community through conferences, seminars, publications and educational services. Founded in 1985, Interop Company is the sponsor of worldwide conferences and expositions, including INTEROP, the premier forum for addressing the interoperability challenges and solutions found in the real world of enterprise computing, from the desktop to the data center. Interop Company is the publisher of *ConneXions—The Interoperability Report*, a monthly technical journal which tracks developments in the computer and communications industry, now in its seventh year of publication.

Interop Company is located at 480 San Antonio Road, Mountain View, California 94040-1219. Interop Europe is headquartered at CNIT BP240, 2, Place de la Défense, 92053 Paris, France.

About Novell, Inc.

Novell, Inc. is an operating systems software company and the developer of network services, specialized and general purpose operating system software products including *NetWare*, *UnixWare* and *DR DOS*. Novell's NetWare network computing products manage and control the sharing of services, data and applications among computer workgroups, departmental networks and across business-wide information systems.

INTEROP is a registered trademark and INTEROPnet is a trademark of Interop Company. NetWorld is a registered trademark of Novell, Inc.



Book Reviews

Effective Management of Local Area Networks: Functions, Instruments, and People, by Kornel Terplan, McGraw-Hill, 1992, ISBN 0-07-063636-2.

No insight

From the title, it would appear that the book is about LAN management, with an emphasis more on process than technology. Unfortunately, as a reviewer it was impossible for me to determine the intended audience. This troubles me quite a bit. The primary input to the book seems to be market data. To be sure there is an attempt to categorize and relate each and every aspect of networking into neat tables and charts in order to form comparisons. Even ignoring some of the rather amusing biases in some of the comparisons (the author doesn't seem to be able to distinguish between marketing hype and real technology), I couldn't really get a sense of any insight from the book. Perhaps I missed some great meaning from the author. Rather immodestly, I think this unlikely.

Network Management: A Practical Perspective, by Allan Leinwand and Karen Fang, Addison-Wesley, 1992, ISBN 0-201-52771-5.

Now here is a book of a different color. In 222 pages, Leinwand and Fang present a concise description of the tasks of a line engineer responsible for network management.

Generic aspects

The book is divided into two major parts: it begins by looking at the generic aspects of network management, as promulgated by the ISO model for systems management: fault management, configuration management, security management, performance management, and accounting management. (The authors' practical perspective is evident in their use of the ISO model only as a tool for categorization, *not* as a paragon of technology.)

While describing each type of management, the book describes the problems encountered and then presents algorithms for tools that can be used to diagnose and potentially correct those problems. This is a welcome departure from books that talk only about models, or only about mechanisms, or only about policies.

SNMP

The second major part of the book discusses how the Internet-standard Network Management Framework can be used for each of the five generic aspects of management. After the usual cursory examination of SNMP and its would-be competitors, the focus is on the Internet-standard MIB and how the objects contained therein can help the network engineer. Finally, the book closes with a discussion on tools that can be built using the Internet-standard Network Management Framework (e.g., MIB compilers).

I really can find only two (small) faults with the book: first, it doesn't have a conclusions section. It just sort of ends. Second, there is a "Series Foreword" and not a real Foreword. The Series Foreword is about as bland as you can get. Considering how much I liked the book, I have to say "shame" to the series editors and complain that *Network Management: A Practical Perspective* deserves better.

DNS and BIND, by Paul Albitz and Cricket Liu, O'Reilly & Associates, 1992, ISBN 1-56592-010-4.

Theory and practice

O'Reilly's "Nutshell" series is known for books which, for a given technology, combine a brief amount of theory with large amounts on useful administration information for an implementation of that technology. And, as the title suggests, *DNS and BIND* is no exception: the technology is the Internet's name-service, the *Domain Name System* (DNS), and the implementation is the *Berkeley Internet Name Domain* (BIND). If you're someone who needs to run BIND (and who really couldn't care less about the details of the DNS), then this is the book for you.

Getting started

The book begins with a motivation as to why a name-service is needed in the Internet community, and how to tell if you need one yourself. Following this, is a very well-thought out section on finding out where in the name-space one should be registered and how to contact the appropriate naming authority.

Next, several chapters deal with installation and configuration of BIND—for a variety of different environments, including those with firewalls. As such, the book contains a step-by-step explanation of how to configure name-service for most readers' environments.

Maintenance


Once you have the name-service configured, the administrator is concerned with maintenance and trouble-shooting. The book provides several chapters on each. In terms of maintenance, in addition to seeing how well things are running, the book also discusses how to deal with growth in a domain and when it is appropriate to delegate authority to a sub-domain. In terms of trouble-shooting, the book goes to great length to discuss several manifestations and how to go about figuring out what is wrong and how to fix it.

My only negative comment is rather minor: the authors offer a one-page comparison of the DNS and X.500, but neglect to explain the primary difference between them (Answer: the DNS is a name-service, X.500 is a directory-service; in a name-service, the primary interrogation service is called *reading*, whilst in a directory-service it is called *searching*.) However, in the grand scheme of things, this criticism is rather minor.

Good reference

A danger in writing a book that is so closely tied to an implementation, is that it might read like part of the documentation set. Have no fear, *DNS and BIND* contains a lot of useful information that you'll never find written down anywhere else. And since it's written in a crisp style, you can pretty much use the book as your primary BIND reference.

—Marshall Rose



INTEROP 93
SPRING

8-12 March 1993 • Washington, D.C. Convention Center

CONNEXIONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President;
Corporation for National Research Initiatives

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute



Printed on recycled paper

Subscribe to CONNEXIONS

U.S./Canada ☐ \$150. for 12 issues/year ☐ \$270. for 24 issues/two years ☐ \$360. for 36 issues/three years

International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

480 San Antonio Road, Suite 100
Mountain View, CA 94040 U.S.A.
415-941-3399 FAX: 415-949-1779

connexions@interop.com

Back issues available upon request \$15./each
Volume discounts available upon request

CONNEXIONS